A Canon Communications LLC Publication

EDN®

VOICE OF THE ENGINEER

SEPT 9
Issue 17/2010
www.edn.com

VENDORS TARGET OSCILLOSCOPE SWEET SPOT

Page 30

DISCOVERING THE LAST UNREALIZED POWER REDUCTION

Page 24

GLASS-BREAKAGE DETECTOR USES ONE MICRO-CONTROLLER

Page 39

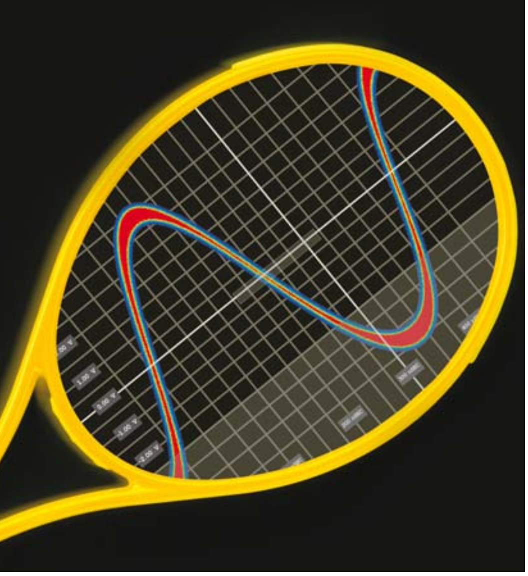

EDN contents
9.9.10

contents 9.9.10

DEPARTMENTS & COLUMNS

EDN online contents www.edn.com

ONLINE ONLY

Check out these Web-exclusive articles:

Designing for EOL with an ASIC
In industries in which product life cycles last a decade or more, the challenges and costs associated with component obsolescence and EOL (end of life) are significant. Engineering teams can address these issues through careful planning and ASIC design.
→www.edn.com/100909toca

The hidden variable: circuit stability as a function of resistor stability
Even close-tolerance instrument components can drift from their specified values. Because of their method of construction, bulk-metal-foil resistors are stable over a wide temperature range.
→www.edn.com/100909tocb

TAKING THE MEASURE

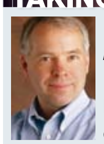

Rick Nelson, editor-in-chief of *EDN* and *Test & Measurement World*, comments on test, globalization, measurement, machine vision, economics, nanotechnology, the engineering profession, and topics of general interest.

Sample a few recent blog entries at the links below:

Getting exercised over technology
→www.edn.com/100909tocc

Is technology more dangerous than darkness and wild animals?
→www.edn.com/100909tocd

The grand challenge of employment
→www.edn.com/100909toce

HOT TOPICS

Hot Topics pages deliver continuously updated, subject-specific links from *EDN* as well as the entire electronics-industry Web.
→www.edn.com/hottopics

EDN® (ISSN#0012-7515) is published semimonthly, 24 times per year, by Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 310/445-4200; FAX 310/445-4299. Periodicals postage paid at Los Angeles, CA, and at additional mailing offices. SUBSCRIPTIONS—Free to qualified subscribers as defined on the subscription card. Rates for nonqualified subscriptions, including all issues: US, $150 one year; $250 two years; $300 three years. Except for special issues where price changes are indicated, single copies are available for $10 US and $15 foreign. For telephone inquiries regarding subscriptions, call 763/746-2792. E-mail: EDN@kmpsgroup.com. CHANGE OF ADDRESS—Notices should be sent promptly to PO Box 47461, Plymouth, MN 55447. Please provide old mailing label as well as new address. Allow two months for change. NOTICE—Every precaution is taken to ensure accuracy of content; however, the publisher cannot accept responsibility for the correctness of the information supplied or advertised or for any opinion expressed herein. POSTMASTER—Send address changes to EDN, PO Box 47461, Plymouth, MN 55447. Canada Post: Publications Mail Agreement 40612608. Return undeliverable Canadian addresses to Pitney Bowes Inc, PO Box 25542, London, ON N6C 6B2.

Volume 55, Number 17 (Printed in USA).

BY RICK NELSON, EDITOR-IN-CHIEF

Learn to apply LEDs—from art to advertising

New LED devices are poised to revolutionize lighting for consumer and other applications. Although they may not be fully competitive with compact fluorescent bulbs for "green"-home lighting, that situation could soon change. LEDs are now finding use in medical-device, automotive, architectural, and signage applications, for which they save power and space and offer the ability to shape light into unlimited colors. This ability lifts them from the realm of the practical into the world of fine art. The work of artist Leo Villareal is evidence of that ability. Villareal works with LEDs to create computer-driven imagery, light sculptures, and site-specific architectural work. You can see his work at the San Jose Museum of Art through Jan 9, 2011.

Not surprisingly, innovators press LEDs into service for marketing as well as art. In June, Huntsman Advanced Materials and the Holst Centre announced the integration of a thin-film-encapsulated flexible OLED (organic light-emitting diode) with the composite material of a LeMans race car's rearview mirror—with the goal of advertising Huntsman's Araldite material to night viewers of the 24-hour race.

From an engineering perspective, to effectively apply HB (high-brightness) LEDs you must understand packaging, control-electronics, and thermal-management issues. You can learn more about these issues at a full-day workshop, "Designing with LEDs," which will take place on Sept 29, 2010, in Chicago. *EDN* and sibling publication *Design News* are sponsors for the event. In four technical-paper tracks you can learn about power management, thermal management, LEDs and solar power, and optics and light measurement. You can also benefit from panel presentations by major LED manufacturers, including Cree, Philips Lumileds, and Osram, which will discuss LED life, reliability, and future products, and see exhibits of the latest in HB LEDs, connectors, packaging technologies, and power-management and thermal-management components.

The workshop is collocated with the Assembly and Automation Technology Expo and five other events at the Donald E Stephens Center in Rosemont, IL. For registration information, go to http://bit.ly/dd1y3o.

When it comes to testing and measuring LEDs, you needn't wait for the workshop. *Test & Measurement World*'s September issue brings you two feature articles on the topic. The cover story recounts Senior Technical Editor Martin Rowe's visit to Luminus Devices, which manufactures LEDs that appear in stage lighting, portable projectors, retail stores, homes, streetlights, and projection TVs. As Rowe writes, lighting designer Kevin Adams chose stage lights that contain Luminus Devices' LEDs when he designed the lighting for Green Day's *American Idiot* musical on Broadway. Rowe explains how the company employs wafer probers, source-measure units, and spectrometers to measure parameters such as forward operating voltage, reverse leakage current, dominant wavelength, and brightness.

In a second September feature in *T&MW*, Bryan C Bolt, technology-development manager for test systems at Cascade Microtech, describes the issues facing test-equipment makers as the HB-LED market experiences a CAGR (compound annual growth rate) projected to exceed 30% over the next several years. Bolt notes that, despite a lack of industry standards and a technology road map, vendors of test equipment must accommodate the range of test configurations they find across the spectrum of manufacturers while controlling the cost of test. He adds that the low-cost mandate suggests that significant customization of test equipment to meet the needs of each LED manufacturer would be impractical and recommends the adoption of purpose-built yet flexible and modular test systems for LED-production test.

It's certain that both design and test engineers have work to do as the applications and markets for HB LEDs expand. EDN

A version of this editorial appeared in the September issue of Test & Measurement World. *You can reach me at richard.nelson@cancom.com.*

EDN. 33 Hayden Avenue, Lexington, MA 02421. **www.edn.com**. **Subscription inquiries:** 1-763-746-2792; EDN@kmpsgroup.com. **Address changes:** Send notice promptly to PO Box 47461, Plymouth, MN 55447. Please provide an old mailing label as well as your new address. Allow two months for the change.

Canon Communications LLC, 11444 W. Olympic Blvd., Los Angeles, CA 90064-1549; 1-310-445-4200; fax: 1-310-445-4299.

Canon Communications LLC

Handheld 8-lb RF-spectrum analyzers boast 43-GHz frequency coverage

Anritsu's new MS272xC series of Spectrum Master handheld RF-spectrum analyzers provide frequency coverage to 43 GHz in units that weigh less than 8 lbs. The series also includes a variety of applications for testing the RF physical layer, making it easier for engineers, field technicians, and monitoring agencies to track over-the-air signals, locate interferers, and detect hidden transmitters. The family comprises five models that cover 9, 13, 20, 32, and 43 GHz.

The units eliminate the need to carry heavy benchtop spectrum analyzers into the field to measure signals at frequencies beyond 20 GHz, such as those in microwave-backhaul applications. To further lighten the load, you can order the new units with a channel scanner and an interference analyzer, with which you can perform all common field measurements, thereby eliminating the need for multiple instruments. The handheld instruments easily incorporate several 3 and 4G (third- and fourth-generation) options to allow measurement of signals that comply with such standards as LTE (long-term evolution), HSPA+ (evolved high-speed packet access), WCDMA (wireless code-division multiple access), EVDO (evolution data-optimized) CDMA, GSM (global system for mobile communication), EDGE (enhanced data rates for global evolution), TD-SCDMA (time-division-synchronous CDMA), HSDPA (high-speed downlink-packet access), and WiMax (worldwide interoperability for microwave access).

The MS2726C takes 27 seconds to sweep a 43-GHz span with a 30-kHz RBW (resolution bandwidth). The units do not sacrifice accuracy for speed, however; they deliver phase noise of –100 dBc/Hz and dynamic range of 101 dB at a 10-kHz offset from a 1-GHz carrier. With a broadband preamplifier to detect small signals, the analyzers also offer high sensitivity. For example, the MS2726C's sensitivity is –159 dBm at 1 GHz and –145 dBm at 43 GHz. An intuitive menu-driven display simplifies all measurements.

An analysis package and Anritsu's Master software tools let you conduct detailed evaluation of measurement data. You can easily identify interference sources using built-in reporting and mapping tools, spectrogram folders, and 3-D spectrograms. These tools eliminate the need for more expensive, larger, heavier benchtop instruments as well as third-party spectrum-monitoring software.

The series uses the field-proven Spectrum Master architecture. A rugged housing withstands the day-to-day operations of field use, and the units' light weight makes them easy to carry when a user is climbing towers. They have a field-replaceable long-life battery and an operating-temperature range of –10 to +55°C. A large, daylight-viewable display eases the viewing of test results in any environment. Display modes include a red night-vision mode, a black-and-white mode, and two full-color modes. US prices start at $15,950.

—by Dan Strassberg

▷**Anritsu Co**, www.us.anritsu.com.

TALKBACK

"Too often, someone will make a 'stupid' comment like 'what if we cycle power?' Where I work, we would all look at him like he is a complete idiot and does not understand real engineering. Of course, we have met the enemy, and he is us."

—Senior design engineer and *EDN* reader Dave McNeely, in *EDN*'s Talkback section, at http://bit.ly/cBBy2c. Add your comments.

The MS272xC series of handheld RF-spectrum analyzers includes five models. This one provides frequency coverage to 43 GHz, which is the widest frequency range in a portable instrument of this type, according to the manufacturer.


HECTRONIC ANOVO CORVALENT NEC SECO
ALT SOFTWARE LYNUXWORKS
MINDTREE QNX SOFTWARE SYSTEMS HECTRONIC VIRTUALLOGIX
REAL-TIME SYSTEMS IBASE TECHNOLOGY MSC VERTRIEBS DIVERSIFIED TECHNOLOGY
EVOC INTELLIGENT TECHNOLOGY BSQUARE
EMERSON NETWORK POWER DFI-ITOX PFU SYSTEMS
NATIONAL INSTRUMENTS ARICENT GE INTELLIGENT PLATFORMS
WIPRO TECHNOLOGIES LSI AVALUE MINDTREE CONCURRENT TECHNOLOGIES NEC
ALT SOFTWARE ADVANTECH IEI TECHNOLOGY VENTURE
PROTECH SYSTEMS TIETO SIGNALING SOLUTIONS RADISYS
WIND RIVER SYSTEMS QNX SOFTWARE SYSTEMS
MACRAIGOR SYSTEMS VIRTUALLOGIX
SENSORY NETWORKS MICRO INDUSTRIES ALT SOFTWARE PROTECH SYSTEMS OKI SEMICONDUCTOR
SECO HECTRONIC MICRO INDUSTRIES
MCS VERTRIEBS ADI ENGINEERING CONTINUOUS COMPUTING
VIRTUALLOGIX PHOENIX TECHNOLOGIES KONTRON XILINX
RED FLAG SOFTWARE AVALUE ARBOR TECHNOLOGY
AXIOMTEK CONGATEC WIPRO TECHNOLOGIES
WIN ENTERPRISES RADISYS PORTWELL
MICROSOFT HECTRONIC PFU SYSTEMS
KONTRON ARIUM AAEON
ARBOR TECHNOLOGY SECO PROTECH SYSTEMS
BSQUARE ANOVO RICOH
PROTECH SYSTEMS ADS-TEC VIRTUALLOGIX
EXAR ADLINK EUROTECH CONCURRENT TECHNOLOGIES ARICENT
TENASYS ANOVO RADISYS WIN ENTERPRISES SUN MICROSYSTEMS
MSC VERTRIEBS RICOH MICROSOFT REAL-TIME SYSTEMS
MSC VERTRIEBS RICOH SENSORY NETWORKS OKI SEMICONDUCTOR
MINDTREE ADVANTECH NETRONOME SYSTEMS DIVERSIFIED TECHNOLOGY
WIN ENTERPRISE CORVALENT ARICENT NEXCOM EUROTECH NORCO
INSYDE SOFTWARE ARICENT ADVANTECH
EMERSON NETWORK POWER NETRONOME SYSTEMS
CURTISS WRIGHT CONTROLS EMBEDDED COMPUTING NATIONAL INSTRUMENTS TENASYS
CONTINUOUS COMPUTING ANOVO LANNER ELECTRONICS CORVALENT
NORCO RED FLAG SOFTWARE
AVALUE EVOC INTELLIGENT TECHNOLOGY EMERSON NETWORK POWER
HEWLETT-PACKARD ALTERA TRENTON TECHNOLOGY
SENSORY NETWORKS WIPRO TECHNOLOGIES
ALTERA AMERICAN MEGATRENDS GREEN HILLS SOFTWARE LYNUXWORKS
NEC EXAR KONTRON BSQUARE ADS-TEC WIND RIVER SYSTEMS NORCO
AVALUE ADI ENGINEERING ARIUM ANOVO

Contact system is first active implantable medical device

Designers of implantable medical-electronic devices for cardiac-rhythm management, neuromodulation, and neurostimulation therapies need ways to connect to those devices. Connection must often occur in a surgical environment, and it's thus important to minimize the time a surgeon needs to make the connections. Addressing this need, Bal Seal Engineering Inc, a provider of custom-engineered sealing, connecting, and conducting products, has announced plans to launch what it claims is the first integrated seal and electrical-contact system for active implantable devices. The new contact system can establish multiple contacts with a single setscrew holding the contacts in place.

Bal Seal is launching an integrated seal and electrical-contact system for active implantable devices. The system can establish multiple contacts with a single setscrew holding the contacts in place.

The company will formally launch the Sygnus implantable contact system next month at the Medical Design and Manufacturing show in Minneapolis. Sygnus combines electrical contacts and isolation seals in a standardized, platform-ready stack configuration. The system pairs the company's Bal Conn electrical-contact technology with pretested silicone seals, resulting in a package that helps medical-device OEMs improve speed to market and eliminate the need for procurement and testing of individual components.

"The upcoming launch of Sygnus is our direct response to the medical-device community's demand for an integrated sealing and connecting system that incorporates an established, ultrareliable contact in a plug-and-play platform that doesn't compromise performance," says Mark Russell, Bal Seal's global medical-electronics account manager. According to Russell, Sygnus builds on technology that's at work in more than a million implantable devices worldwide.

The number of contacts and seals in the Sygnus system will be configurable to meet application and industry requirements, such as the IS4 standard for cardiac-health-management devices. The company's trademarked Canted-coil spring-contact design offers multipoint conductivity and low insertion force, and it compensates for any surface irregularity. The system comprises a series of electrical contacts, including platinum-iridium Canted-coil springs in metal housings made from MP35N, and implantable-grade silicone seals, which provide dielectric isolation. The combined stack is resistant to long-term material adhesion and potential fatigue due to multiple lead-insertion cycles. The company plans to subject the stack and package contacts to force and resistance tests and to package the sealing components to comply with critical standards for "clean" packaging. **—by Rick Nelson**

▷**Bal Seal**, http://balseal.com.

DILBERT By Scott Adams

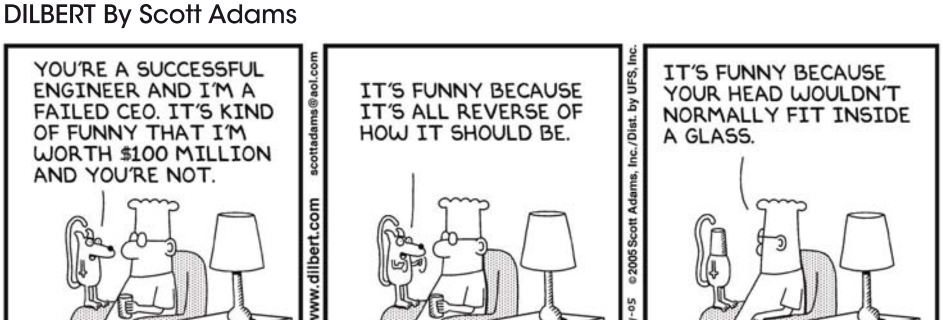

09.09.10

SUPERCAPACITOR CHARGERS TOUT AUTOMATIC CELL BALANCING

Linear Technology's new LTC3625 and LTC3625-1 two-cell supercapacitor chargers address high-peak-power, data-backup, and "dying-gasp" needs in portable-system and data-storage applications. The devices' switch-mode topology includes an internal buck converter between the input voltage and the series capacitor's midpoint voltage to regulate the voltage on the bottom capacitor. It also includes an internal boost converter between the midpoint voltage and the output voltage to regulate the voltage across the top capacitor.

The devices can charge two supercapacitors in series from a 2.7 to 5.5V current-limited supply to a pin-selectable output voltage of 4.8 or 5.3V for the LTC3625 and 4 or 4.5V for the LTC3625-1. They also feature automatic cell-balancing, which maintains approximately equal voltages across both cells, eliminating the need for balancing resistors.

Other features include overtemperature and reverse-current protection and overcurrent limiting. The LTC3625 and LTC3625-1 come in 3×4-mm DFN packages and operate over –40 to +125°C. Prices start at $3 and $3.45 for E- and I-grade versions.

—by Fran Granville

▶**Linear Technology Corp, www.linear.com/3625.**

LabView 2010, VNA, data-acquisition, and I/O products debut at NIWeek

During last month's NIWeek, in Austin, TX, National Instruments introduced LabView 2010, a 6-GHz VNA (vector-network analyzer), data-acquisition modules, and a RIO (reconfigurable-I/O) chassis. LabView 2010 features a rewritten compiler that increases execution speed by an average of 20%, according to the company, with execution of some functions, such as parallel for loops, occurring nearly 200% faster. With LabView 2010, NI has improved the compiler data flow's intermediate representation and added the low-level virtual-machine open-source compiler infrastructure to the software's compiler flow to accelerate code execution.

NI's new 6-GHz, two-port NI PXIe (Peripheral Component Interconnect Express Extensions for Instrumentation)-5630 VNA module targets automated design validation and production test. Its two-slot PXI footprint enables test engineers to incorporate vector-network analysis into their test systems without the added cost and large footprint of traditional bench-top VNAs.

LabView 2010 features a rewritten compiler that increases execution speed by an average of 20%, with execution of some functions, such as parallel for loops, occurring nearly 200% faster.

The PXIe-5630 features automatic precision calibration, full vector analysis on both ports, reference-plane extensions, and a LabView API (application-programming interface) that supports parallel test. The VNA features a frequency range of 10 MHz to 6 GHz, a dynamic range of greater than 110 dB, and sweep speeds of less than 400 µsec/point over 3201 points. Engineers can combine as many as eight modules in one PXI chassis and perform multisite RF test. The base price for the VNA is $25,999.

NI also introduced the Ethernet-based NI CompactDAQ modular data-acquisition system, which includes new cDAQ-9188 chassis that hold eight I/O modules for measuring as many as 256 channels of electrical, physical, mechanical, or acoustic signals in a rugged, 25×9×9-cm form factor. The base price for the chassis is $1399. For USB (Universal Serial Bus) applications, NI X Series multifunction data-acquisition devices integrate high-performance analog measurement and control channels, digital I/O, and counter/timers. USB X Series data-acquisition devices include as many as 32 analog inputs, four analog outputs, 48 digital I/O lines, and four counters. The devices' sampling rates range from 500k samples/sec multiplexed for analog inputs to 2M samples/sec/channel for simultaneously sampling analog inputs. The base price is $1149.

The new 9157 and 9159 MXIe (Multiplatform Extensions for Instrumentation Express) RIO chassis and 9148 Ethernet-RIO chassis complement the NI 9144 EtherCAT (Ethernet-control-automation-technology) chassis. New C Series products extend the company's offering of high-channel-count chassis to a variety of buses. The base prices for MXIe RIO and Ethernet RIO chassis are $4499 and $999, respectively.—**by Rick Nelson**

▷**National Instruments**, www.ni.com.

MICRON UNVEILS SOLID-STATE DRIVES FOR ENTERPRISES

09.09.10

Micron Technology recently introduced its P300 solid-state drives, enterprise-tailored variants of the C300 drives and the latest iteration of Micron's multigenerational stab at this market. Last December, the company unveiled the consumer-tuned C300 line, which it based on 34-nm MLC (multilevel-cell) NAND-flash memories from the fab it shares with partner Intel (www.intel.com).

Micron has had a few firmware glitches with the C300 series, leading to access lockouts, data corruption, and other problems. As of mid-May, however, things seemed to finally be stable with the C300 series, giving Micron the green light to roll out the P300 siblings in 50-, 100-, and 200-Gbyte flavors. The company based the devices on SLC (single-level-cell) flash memories, which provide inherently higher data integrity at a given block-erase-cycle count than their lower-cost-per-bit MLC counterparts. SLC-NAND-flash memory also delivers higher write speeds than MLC alternatives. And, like the C300 products, P300 drives harness a 6-Gbps SATA (serial-advanced-technology-attachment) system interface.

Micron also uses ECC (error-correcting-code)-enhanced DRAM in the P300 series, which the company claims will sell for a competitive price of less than $10/Gbyte. About that system interface, however, Micron hasn't yet equipped its P drives with SAS (serial-attached SCSI). SCSI (small-computer-system-interface) hooks are deeply embedded in enterprise code, and a lack of corresponding hardware support can be a deal breaker. Dean Klein, vice president of memory-system development at the company, claims that SATA support is sufficient for many enterprise applications. Still, Micron competitors Toshiba (www.toshiba.com), Samsung (www.samsung.com), and Seagate (www.seagate.com) are gearing up for SAS, both internally and in partnerships.

Micron may instead choose an alternative next-generation enterprise-interface path. This is the tack that Fusion-io (www.fusionio.com), for example, takes in harnessing PCIe (Peripheral Component Interconnect Express) to link its SLC- and MLC-based flash-memory boards to systems.—by Brian Dipert

▶**Micron, www.micron.com.**

Process aims to make solar power cheap enough to compete with oil

Stanford University engineers claim to have discovered an energy-harvesting process that could surpass the efficiency of today's photovoltaic and thermal-conversion technologies. The university's process coats a piece of semiconducting material with a thin layer of cesium, making the material able to use both the light and the heat of the sun to generate electricity.

The process works well at high temperatures, unlike the photovoltaic technology that solar panels currently use, which become less efficient as the temperature rises.

The PETE (photon-enhanced-thermionic-emission) process could make solar-power production more than twice as efficient as current methods and potentially cheap enough to compete with oil, according to the researchers. "This is really a conceptual breakthrough, a new energy-conversion process, not just a new material or a slightly different tweak," says Nick Melosh, an assistant professor of materials science and engineering at the school and the leader of the research group. "It is actually something fundamentally different about how you can harvest energy."

Although high temperatures are necessary to power heat-based conversion systems, solar-cell efficiency rapidly decreases at higher temperatures. According to the university, heat from unused sunlight and inefficiencies in silicon cells account for a loss of more than 50% of the initial solar energy reaching the cell. The Stanford team's work focused on wedding thermal- and solar-cell-conversion technologies.

Although most silicon solar cells are inert by the time the temperature reaches 100°C, the PETE system hits peak efficiency at temperatures higher than 200°C, making it useful for large-scale solar farms in a desert, for example. "We've demonstrated a new physical process that is not based on standard photovoltaic mechanisms but can give you a photovoltaic-like response at very high temperatures," Melosh says. "It works better at higher temperatures."

According to a paper describing the research, the researchers used a gallium-nitride semiconductor to test PETE because it withstands high temperatures (**Reference 1**). Other materials, such as gallium arsenide, could reach as much as 60% efficiency, the researchers estimate. The materials to build a device to make the PETE process work are inexpensive and easily available, according to the researchers, making the power that comes from it affordable. The team wants to design the devices so users could easily bolt them onto current systems, thereby making conversion relatively inexpensive. "The material cost … is not really an issue, unlike the way it is for large solar panels of silicon," says Melosh.

The Global Climate and Energy Project at Stanford and the Stanford Institute for Materials and Energy Science, a joint venture of Stanford and SLAC (Stanford Linear Accelerator Center) National Accelerator Laboratory, provided funding for the research, with additional support from the Department of Energy and the Defense Advanced Research Projects Agency.

—by Suzanne Deffree

REFERENCE

1. Schwede, Jared W, et al, "Photon-enhanced thermionic emission for solar concentrator systems," *Nature Materials*, Volume 9, Aug 1, 2010, pg 762, http://bit.ly/aCBJiL.

▷**Stanford University**, www.stanford.edu.

Circuits claim 1000-times efficiencies in cost, power, and size over today's digital computing

09.09.10

Lyric Semiconductor has launched a probability-processing technology, which it believes will in the future offer 1000 times more efficiency in cost, power, and size than today's digital computing. The new technology enables many applications that now require 1000 conventional processors to run in just one of Lyric's processors. Probability processing computes likelihoods, or odds. Its logic-gate circuit uses transistors as dimmer switches instead of as on/off switches. Lyric's circuits can accept inputs and calculate outputs between zero and one, directly representing probabilities, according to the company.

Lyric's first commercialized application of the probability-processing technology, the LEC (Lyric Error Correction) for flash memory, offers 30-times smaller cores and ASICs and a 12-times decrease in power consumption at higher throughput than digital approaches.

Lyric's LEC for flash memory offers 30-times-smaller cores and ASICs and a 12-times decrease in power consumption at higher throughput than digital approaches.

"After a decade of development, we have no shortage of opportunities for our probability-processing technology, but we are currently focused on a modest list of both short- and long-term applications that will see enormous gains in performance," says Ben Vigoda, PhD, the company's chief executive officer and co-founder.

Lyric ultimately plans to develop the GP5 (general-purpose programmable probability-processing platform), which will calculate probabilities for all types of applications—from Web searches to genome sequencing—and could allow for performance gains over current digital x86-based systems. GP5, which should become available for sampling in 2013, will run code written in the company's PSBL (probability-synthesis-to-Bayesian-logic) language.—by Suzanne Deffree

▶**Lyric Semiconductor, www.lyricsemiconductor.com.**

AVAGO TECHNOLOGIES

New Smaller and Faster Optical Isolation Amplifiers Feature 0.5% Gain Accuracy and 1140 Vpeak Working Voltage

Introduction

Many analog designers are familiar with differential instrumentation amplifiers but these will not provide the insulation voltages to withstand high transient voltages safely, nor the isolation to protect sensitive low voltage control electronics from high voltage switching circuits found in power conversion applications. With a -40° to 105°C operating temperature range, Avago's miniature ACPL-C79x Precision Isolation Amplifiers target industrial automation and instrumentation, renewable energy, and HVAC markets.

Based on Avago's proprietary optical isolation technology, sigma-delta analog-to-digital converters and chopper stabilized amplifiers, the ACPL-C79x isolation amplifiers are used for motor phase and rail current sensing, servo motor drive, switching power supply feedback isolation, DC link voltage monitoring, inverter current sensing and switching power supply feedback isolation. The ACPL-C79x high common-mode transient immunity of 15 kV/µs provides the ruggedness and stability needed to accurately monitor current in high-noise motor control environments.

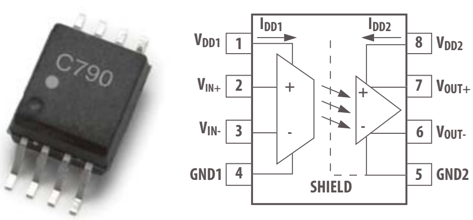

Figure 1. ACPL-C79X Package

As shown in Figure 1, the isolation amplifiers are fully differential, input and output, with a gain accuracy of ±0.5% (ACPL-C79B), ±1% (ACPL-C79A), and ±3% (ACPL-C790). Operating from a single 5 V supply, the isolation amplifier series features an excellent nonlinearity of 0.05% and a SNR of 60 dB. With a 200 kHz bandwidth and 1.6 µs response time, the ACPL-C79x captures transients during short circuit and overload conditions. The stretched SO-8 package has a footprint 30% smaller than the standard DIP-8 package. When mounted on a PCB, it occupies a space that is a fraction of that for a Hall Effect or transformer based isolation amplifier.

Key ACPL-C79x Key Features

- Fully Differential Isolation Amplifier
- ±0.5% High Gain Accuracy (ACPL-C79B)
- –50 ppm/°C Low Gain Drift
- 0.6 mV Input Offset Voltage
- Excellent 0.05% Linearity
- 60 dB SNR
- 200 kHz Wide Bandwidth
- 3 V to 5.5 V Wide Supply Range for Output Side
- -40°C to +105°C Operating Temperature Range
- Advanced Sigma-Delta (Σ-Δ) A/D Converter Technology
- 15 kV/µs Common-Mode Transient Immunity
- Safety and Regulatory Approvals (pending):
  - IEC/EN/DIN EN 60747-5-5: 1140 Vpeak working insulation voltage
  - UL 1577: 5000 Vrms/1min double protection rating
  - CSA: Component Acceptance Notice #5

Motor Drive Application Example

In a typical motor drive application, shown in Figure 2, currents through a small value current sense resistor cause a voltage drop that is sensed by the ACPL-C79x and a differential output voltage, proportional to the current, is created on the output side of the isolation barrier. A floating power supply (which in many applications could be the same supply that is used to drive the high-side power transistor) is regulated to 5 V using a simple three terminal voltage regulator (U1). The voltage from the current sensing resistor, or shunt (RSENSE), is applied to the input of the ACPL-79x through an RC anti-aliasing filter (R5 and C3). And finally, the differential output of the isolation amplifier is converted to a ground referenced single-ended output voltage with a simple differential amplifier circuit (U3).

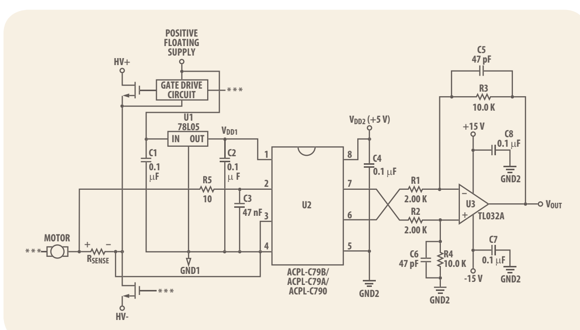

Figure 2. Typical motor current sense circuit

Although the application circuit is relatively simple, a few recommendations should be followed to ensure optimal performance.

Shunt Resistor Selection

A real-world motor current sense resistor calculation will show what must be considered in selecting the current sense resistor. First determine how much current the resistor will be sensing. The graph in Figure 3 shows the RMS current in each phase of a three phase induction motor as a function of average motor output horsepower and motor drive supply voltage. The maximum value of the sense resistor is determined by the current being measured and the maximum recommended input voltage of the isolation amplifier. For example, if a motor has a maximum RMS current of 10 A and can experience up to 50% overloads during normal operation, then the peak current is 21.1 A (=10 x 1.414 x 1.5). With a maximum amplifier input voltage of 200 mV, the maximum sense resistance would be about 10 mΩ. The maximum average power dissipation in the sense resistor, which is about 1 W in this example, should also be checked by multiplying the sense resistance times the square of the maximum RMS current. If the power dissipation is excessive, a lower resistance value can be used.

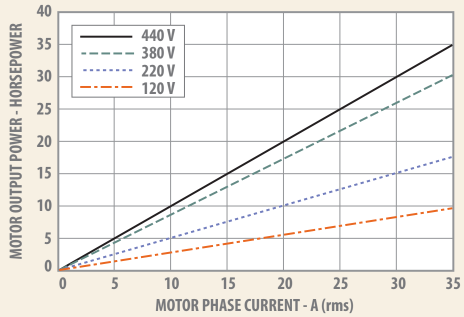

Figure 3. Horsepower vs. motor phase current/voltage

Differential Input Connection

In Figure 2, the isolation amplifier is connected in a single-ended input mode. However, given the fully differential input structure, a differential input connection, shown in Figure 4, can be used for better performance. Any noise induced on one pin will be coupled to the other pin by the capacitor C and creates only common mode noise which is rejected by the ACPL-C79x.

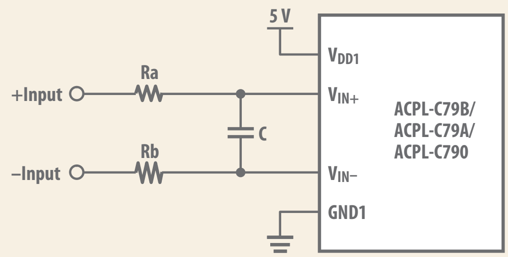

Figure 4. ACPL-C79x differential input connection

Voltage Sensing

The ACPL-C79B/C79A/C790 can also be used to isolate signals with amplitudes larger than its recommended input range with the use of a resistive voltage divider at its input. The only restrictions are that the impedance of the divider be relatively small (less than 1 kΩ) so that the input resistance (22 kΩ) and input bias current (0.1 µA) do not affect the accuracy of the measurement. An input bypass capacitor is still required, although the 10 Ω series damping resistor is not (the resistance of the voltage divider provides the same function). The low-pass filter formed by the divider resistance and the input bypass capacitor may limit the achievable bandwidth.

Evaluate the ACPL-C79X

The ACPL-C79X evaluation board demonstrates the high linearity and low-offset capability of the ACPL-C79B/C79A/C790. It allows a designer to easily test the performance of the isolation amplifier in an actual application under real-life operating conditions. A surface mount shunt resistor is provided along with the board.

Figure 5. ACPL-C79x evaluation board

Summary

Avago's three new Miniature Precision Isolation Amplifiers, made possible by proprietary optical isolation technology, offer increased accuracy, speed, bandwidth and insulation ratings. Additional information is available at www.avagotech.com

BY HOWARD JOHNSON, PhD

Linearity

My good friend Chris "Breathe" Frue is a talented musician, a trained audio engineer, and an excellent conversationalist. He asked recently, "What is the meaning of linearity, and why should I care?"

I took a long puff on my pipe and answered slowly, "Well, linearity is one of two properties essential for good signal fidelity—audio or otherwise. The other property is time invariance. A linear, time-invariant system responds equally well to loud and to soft inputs, whether composed of one sound or many."

"You are just waving your hands," Breathe said. "I don't buy it. There must be some more-concrete definition."

"There is," I replied. "It's tricky to state the whole thing, so I'll begin with a necessary condition, meaning that every linear system must at least do this task. The condition is called scaling. Scaling means that, if you turn up the volume on the system input, the system response scales proportionately. Your guitar amplifier, for example, has the property of scaling." (Breathe plays a fine old arch-top jazz guitar. He uses a Mackie mixer driving a linear studio-quality monitor to produce a clean sound. He doesn't need distortion because his technique is impeccable.)

"I don't believe that," Breathe said. "Look, if I tweak the volume knob on my guitar to 5, it sounds one way. If I turn it up to 10, the club manager comes over and tells me to turn it down. So the response is totally different in those two cases."

"Yes," I answered. "And your speaker probably distorts at the high setting, too, so that won't be the same either, but what I'm saying is that, if you keep the volume in a reasonable range, then scaling works."

"Does the term 'reasonable range' include a setting of zero?" he asked.

"Of course, zero is a perfectly valid input signal for any linear system," I said. "The output would be zero."

"But it's not," said Breathe. "Even when I set my guitar to zero, a little hiss always comes out of the speakers. So the amp is not, according to your definition, linear for either large-scale or small-scale inputs."

At this point, I realized that, through earlier such conversations, I had already taught Breathe far too much about electrical engineering. His questions were becoming dangerous. My next columns will lay out for him, in a methodical but simple way, the whole concept of linear-time-invariant behavior so he can understand its importance, not only as a tool for modeling but as an ideal standard of behavior against which you can measure circuit performance.

ON ANOTHER TOPIC, a recent article (**Reference 1**) generated a lot of reader responses. Here are some of the more oft-repeated ideas. First, the precise values for the 10% resistor scale (10, 12, 15, 18, 22, 27, 33, 39, 47, 56, 68, and 82) nearly fit an exponential scale. The steps are adjusted so that the tolerance bands in most cases overlap. For example, the nominal value of 68Ω–10% gives 61.2, slightly smaller than 56+10%, which equals 61.6. Only the gaps from 12 to 15Ω and 18 to 22Ω violate this rule. Because most

I had taught him far too much about electrical engineering. His questions were becoming dangerous.

of the bands overlap, almost any resistor you manufacture fits into some tolerance band somewhere on the scale. Few parts go to waste, and manufacturers love this fact. The other tolerance scales—20, 5, and 2% and so on—have similar overlapping properties.

Next, if you file the side of a carbon-composition resistor, notching through its outer coating into the bulk carbon layer, you can raise its resistance. This approach makes every resistor a "variable resistor." A drop of lacquer reseals the outer coating. Don't file too far!

Finally, the value of a carbon-composition resistor drifts with temperature and with age. If you want long-term stability, you must prebake your resistors. **EDN**

REFERENCE

1 Johnson, Howard, PhD, "7% solution," *EDN*, June 10, 2010, pg 22, www.edn.com/article/509250-7_solution.php.

Howard Johnson, PhD, of Signal Consulting, frequently conducts technical workshops for digital engineers at Oxford University and other sites worldwide. Visit his Web site at www.sigcon.com.

BY PALLAB CHATTERJEE, CONTRIBUTING TECHNICAL EDITOR

Software for nanodesign

Nanotechnology is a diverse area that primarily involves specialty-material properties and effects that are prominent only at nanoscale dimensions. The focus of the designs in these areas is as diverse as the materials. As a result, the CAD (computer-aided-design)-software community has not come up with a unified approach. The market has instead chosen to apply generalized mathematical and materials-based software programs to help with the task.

The mathematical tools, including Wolfram Research's (www.wolfram.com) Mathematica, The MathWorks' (www.mathworks.com) Matlab and Simulink, and PTC's (www.ptc.com) MathCAD, determine the materials' empirical properties and the particles' actions. These general-purpose mathematic solvers have for decades been the mainstays of physics, chemistry, and engineering groups as the primary methods for addressing new technologies and techniques by solving the primarily nonlinear or piecewise-linear fundamental equations governing the design or molecular dynamics.

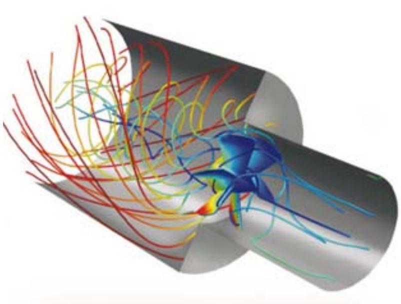

Figure 1 Comsol's Multiphysics tool lets you model interaction analysis, such as airflow, independently of geometries for a device.

Physics-based CAD and computational tools, targeting advanced visualization, show the typical results for specific types of analysis. These tools include Comsol's (www.comsol.com) Multiphysics, Silvaco's (www.silvaco.com) TCAD (technology-CAD), and Verseon's (www.verseon.com) DDA (drug-design automation).

Comsol's Multiphysics dominates the market due to its specialized computational engines and ability to interact with most 3-D-design tools and with mathematical-modeling products, such as Matlab. Multiphysics has modules for ac/dc, RF, MEMS (microelectromechanical systems), plasma, structural mechanics, acoustics, heat transfer, chemical-reaction engineering, batteries and fuel cells, earth science, and CFD (computational fluid dynamics). You can model interaction analysis, such as airflow, independently of geometries for the devices (**Figure 1**). It is an equation-based tool set, so, if you use the proper information for nanoscale and microscopic structures, you can analyze them in an equivalent way to how you would analyze macroscopic structures.

Specialized TCAD tools allow you to create and model virtual processes and to characterize devices. **Figure 2** shows the Silvaco TCAD tool's organic modeling of OLED (organic light-emitting-diode) displays for a photon excitation. These single-particle events are the keys to the development of nanotechnology structures, and they are difficult to reproduce in a laboratory.

Pharmaceutical and biomedical nanotechnology is now focusing on molecular dynamics and chemical reactions that form the basis of the medication- and drug-creation industry. Verseon's DDA product blazes a trail for a new direction in the dynamic simulation of drugs. It allows for the virtual creation, simulation of interaction, and quantitative analysis of the combination of known molecules into new structures. The traditional method involves processing the molecules into structures, separating the elements of interest, growing sufficient quantities of these elements for testing, and then performing laboratory-grade analysis to see the elements' interaction with other compounds. Computers simplify these tasks, dramatically reducing the time it takes to perform them, and allow engineers to analyze the sensitivities to single or low numbers of interacting particles. This capability also reduces wasted time on synthesis paths that show unpromising results.

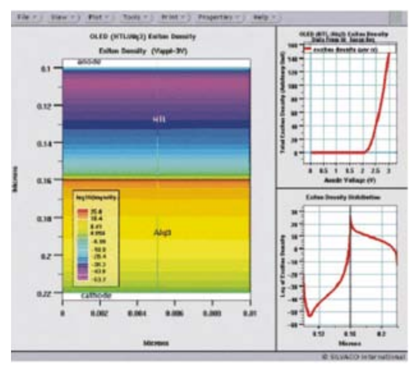

Figure 2 Silvaco's TCAD tool performs organic modeling of OLED displays for a photon excitation.

At the nanoscale, computer-simulation aspects are essential due to the lack of granularity and access to the devices available in the macroscale. EDN

Pallab Chatterjee is vice chairman of the IEEE San Francisco Bay Area Nanotechnology Council. You can reach him at pallabc@siliconmap.net.

Single Box Solution at Low Cost Per Channel

High Speed Multi-Channel BERT

Multi-Channel
4-tap De-emphasis

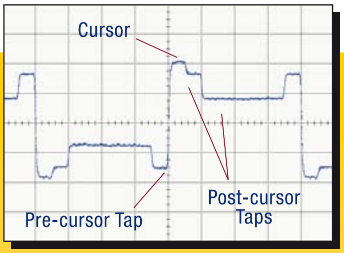

Multi-Channel Asynchronous Traffic Generation for Real Time Crosstalk Measurement

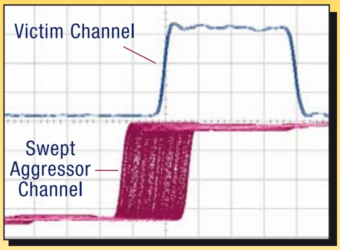

Multi-Channel Jitter Tolerance

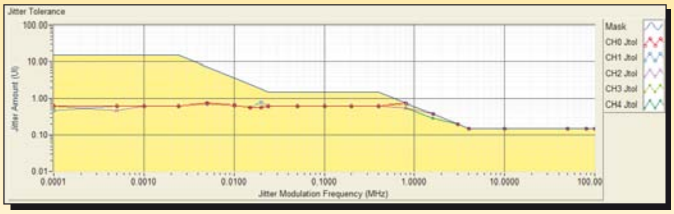

Key Features

- Easy to use GUI for multi-system integration
- Modular architecture supports 1 to 5 remote generator or detector heads
- Integrated 4-tap de-emphasis
- GPIB/USB controllable pattern generator and error detector for automated measurements
- Transparent jitter pass-through
- Swept aggressor amplitude and phase delay for crosstalk characterization
- Precision phase alignment of clock or data

Applications

- Crosstalk characterization
- Multi-Channel Jitter Tolerance testing on high speed multi-lane systems
- Parallel Bit Error Rate testing for higher manufacturing throughput
- Multi-channel BER tester for CFP (10x10G) and QSFP (4x10G) standards
- Testing backplanes and connector performance in the presence of multiple aggressors
- Testing 10G Ethernet 10G BASE-KR
- Forwarded clock receiver characterization

CENTELLAX

3843 Brickway Blvd. • Suite 100 • Santa Rosa • CA 95403 • USA
Ph 707.568.5900 | **Fax** 707.568.7647 | sales@centellax.com
Toll Free 866.522.6888 | www.centellax.com

For additional details, application notes and assembly diagrams, please visit our website: www.centellax.com

DISCOVERING THE LAST UNREALIZED POWER REDUCTION

POWER-OPTIMIZED ARCHITECTURES HELP ENGINEERS DESIGNING CHIPS WITH BLOCKS THAT CAN POWER DOWN OR OPERATE AT REDUCED FREQUENCIES AND VOLTAGES.

BY JAY CHIANG • SYNOPSYS

Power has become one of the most important design criteria for almost all design projects, and the industry, in response, has invested a lot of effort to address this challenge. Consequently, we have seen a plethora of low-power design techniques and new technologies emerge. Some of these techniques are relatively easy to adopt. For example, clock gating and multiple-threshold-voltage cells have become mainstream design practices because they are effective. In addition, EDA tools can automate their implementation. Some techniques, on the other hand, require more planning. For example, design engineers can group SOC (system-on-chip) circuits into multiple blocks so that they can power down some blocks or operate them at reduced frequencies or voltages when operating conditions allow it. Although these more advanced techniques take more deliberate effort to implement, design engineers are increasingly employing them to meet the more stringent power requirements in next-generation chips.

When applying low-power design techniques, design engineers typically concentrate on only the few modules, such as embedded processors and on-chip memories, that consume more power than the other blocks. Although this focus is necessary, it is incomplete. Engineers may often overlook the fact that many low-power-consuming blocks frequently have a greater impact on energy consumption than their power-consumption number suggests. If you correctly plan a chip's power-management strategy, the power-consumption profile and energy-consumption profile should not correlate closely. You should keep the active period of the high-power-consuming modules as short as possible. The modules that remain powered for a long time should not consume too much power. Even though these modules consume less power than other blocks, they consume a higher propor-

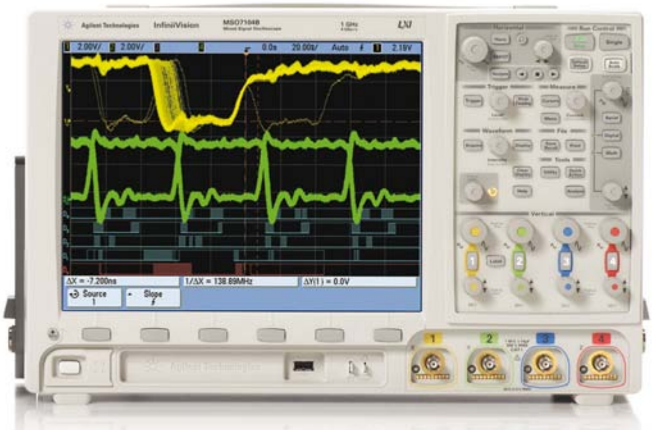

Agilent InfiniiVision 7000B Series (MSO and DSO models) Bandwidth: 100 MHz - 1 GHz

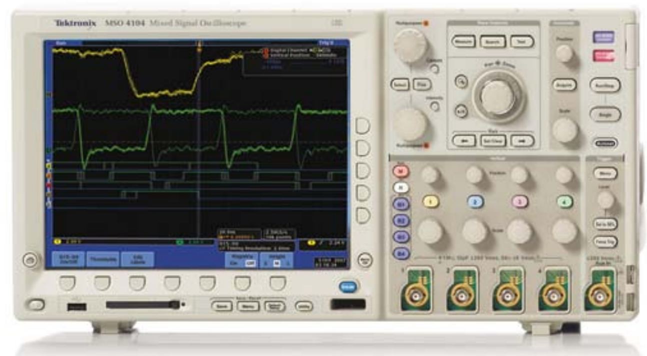

Tektronix 4000 Series

40% bigger screen.

475,000 times faster...Yes, really.

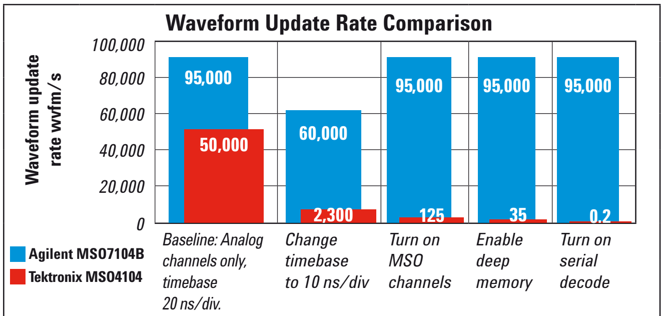

Notes:
- Measurements taken on same signal using Agilent MSO7104B and Tektronix MSO4104 .
- Memory depth = display window X sample rate with up to 8 Mpts on Agilent.
- Tektronix measurements taken with version 2.13 firmware.
- Screen images are actual screen captures, and scopes are shown to scale.

A scope's real world performance depends on more than just bandwidth and sample rates. It takes an ultra fast update rate to capture infrequent events and display the most accurate signal detail. That's where Agilent 7000B Series scopes really stand out. Thanks to a proprietary chipset, you get faster update rates not only in analog mode, but in every mode. That's why you see greater signal detail. That's Agilent.

Agilent and our Distributor Network

Right Instrument.
Right Expertise.
Delivered Right Now.

RS ELECTRONICS®

888-852-7234
www.rselectronics.com

Are you using the best scope?
Take the 5-minute Scope Challenge.
Visit www.agilent.com/find/faster

Agilent Technologies

tion of energy once you factor in their extended active time.

Consider a hypothetical cellular-phone design. Under typical usage, the cellular phone is mostly in standby mode. During standby, most circuits, except the wireless receiver or receivers, are off. Although standby mode consumes only a fraction of the power that the other modes consume, it still consumes 36% of the total energy, after factoring in the active period. In other words, it pays dividends to aggressively reduce power for circuits that are active in the standby mode because it can lead to significant savings in battery life (**Table 1**).

Such opportunities for energy reduction exist in most SOCs. In general, if the chip has multiple power domains, it has multiple power modes. If you identify the power modes that are most active, you can isolate the circuits that have higher impact on the chip's energy consumption, and you can more aggressively pursue power reduction in these focused areas to reduce the overall energy footprint of the chip.

Analysis of these circuits in further detail uncovers some interesting characteristics. These modules must remain on for extended periods because they perform essential functions for the chip in that operating mode. They are often continuously calculating data or processing signals. In addition to the cellular-phone example, other circuits, such as audio or video processors in playback or talk mode and signal-processing blocks, such as equalizer, modulation, or cryptology units, in wireless and networking applications, have more datapath content than control logic and can benefit considerably from low-power techniques.

If you consider the technology horizon, a new generation of connected devices aiming to deliver better user experiences and higher data rates is driving many new design starts. Consequently, these new projects will demand higher

AT A GLANCE

- Design engineers are increasingly employing advanced techniques to meet the more stringent power requirements of next-generation chips.
- It pays dividends to aggressively reduce power for circuits that are active in standby mode because it can lead to significant savings in battery life.
- Power gating isn't feasible for circuits that must continuously remain active, so the only choice is to make the circuit intrinsically low power.
- Traditionally, datapath generators produce the most area-economic architectures that still meet the timing constraints.
- Because power is a physical-domain characteristic, your standard-cell library can affect the power-optimization result.

audio quality, higher video resolution, more pixel support, more complex signal processing, faster data rates, and so forth. Increases in the size and complexity of the signal-processing blocks in turn lead to a higher energy footprint in the new designs. The impact of this design complexity requires design engineers to more closely manage the power consumption for these blocks.

LOW-POWER DATAPATHS

Power gating isn't feasible for circuits that must continuously remain active, so the only choice is to make the circuit intrinsically low power. The first step is to lower the voltage, the operating frequency, or both without missing the performance target. However, slower clock frequencies mean deeper logic levels, and these circuits usually include more datapath logic than control logic. Datapath logic is notoriously prone to glitches—unwanted transitions that settle before the next clock edge—and switching because any spurious transitions propagate downstream and ripple throughout the entire datapath tree. Although glitches pose no functional issues, these transitions still consume power.

It is critical to avoid increasing power in other areas while reducing it in one area. Making this power-reduction approach more effective requires more balanced, shallower architectures that can limit the propagation of the transitions. Although most EDA tools do an adequate job producing timing- and area-optimized architectures that designers later optimize for power at the gate level, they are less effective in considering the power consequence of architectural selections upfront.

Some design engineers try various means of writing power-optimized architectures into RTL (register-transfer-level) code to save power. However, most low-power architectural-RTL coding focuses on reducing area, based on the assumption that using fewer cells equates to less power consumption. For example, some design engineers in networking and multimedia applications truncate the LSBs (least-significant bits) of the data when precision is not critical.

Although this technique is useful, you must understand the details of how to implement it. Datapaths differ from other logic circuits in that they perform computer arithmetic that generates carries and sums, requiring carry-propagating adders to add together the carry and sum to produce a binary number. For RTL coded at a high level, EDA tools usually can generate datapath architectures, keeping all the numbers in redundant format—annotating the value of the number with both carry and sum—until the last level of the output.

If you code the datapath at a lower level, you might turn to coding practices that divide a larger datapath block into several small ones, forcing the RTL-synthesis tools to insert carry-propagation adders into the final stage of every smaller block (**Figure 1a**), hence increasing area and delay.

TABLE 1 SAMPLE POWER MODES AND ENERGY CONSUMPTION

Power mode	Power consumption (mW)	Time budgeted in mode (%)	Energy-consumption profile
Standby	40	90	36
Audio	400	3	12
Phone	500	5	25
Video	1200	2	24

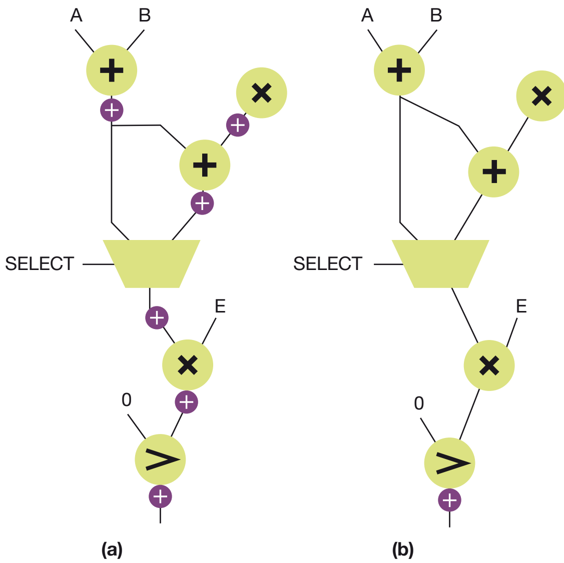

Figure 1 Multiple carry-propagation adders for every fragmented datapath block (a) increase power consumption. Merging datapath blocks and providing one carry-propagation adder for each merged block (b) avoids unnecessary binary conversions.

The resulting increased area sometimes offsets the entire power gain from the LSB truncation. For optimal results, you must consider RTL-coding practices that allow the merging of datapath blocks to avoid unnecessary binary conversions (**Figure 1b**).

Some design engineers also try to code isolation logic in front of the datapath logic so that they can suppress the switching and transition of the datapath tree until there is valid data. Depending on the input-data profile and how frequently the data is valid, this approach could save significant dynamic power. The concept, operand isolation, is similar to clock gating, except that it takes place on the datapath instead of the clock paths (**Figure 2**). The concept, also known as data gating or datapath gating, is appealing, but it is sometimes difficult to implement in practice. Unlike clock gating, adding isolation logic to datapaths increases the path delay. This timing overhead can make it tricky to close timing. Some RTL-synthesis tools can automatically insert the isolation logic; however, engineers do not widely use the feature because it degrades timing.

AN ALTERNATIVE APPROACH

Datapath generators traditionally produce the most area-economic architectures that still meet the timing constraints. Engineers then optimize the generated designs for power at the gate level. At this level, the scale of optimization involves only a few gates. The flows don't provide power-optimized architectures, so some designers manually code them in low-level RTL, which can hinder datapath optimization and degrade the quality of results.

To improve this situation, the first step is to understand what kind of datapath architectures consume less power so that you can use the knowledge to create more low-power architectures. Second, you should characterize the power costs of the datapath structures at a high level so that you can fully consider the power consequences when making architectural decisions.

Examples include the power-stingy architectures of the Synopsys (www.synopsys.com) DesignWare minPower components. These low-power datapath architectures are flatter, shallower, and more balanced than traditional architectures to produce fewer spurious transitions. When these unwanted transitions occur, datapath structures with smarter cell selections can limit their propagation. For example, instead of using common XOR-based datapath cells, such as full adders or XOR-based booth encoders, the manpower components employ architectures that favor more AND or NAND cells so that fewer transitions ripple throughout the datapath tree.

Integrating these power-friendly architectures yields some advantages. Aside from being easier to use, these architectures allow designers to capture power-saving opportunities that are hard to realize with a manual approach. Because power consumption depends on operating conditions, it is not enough to consider the circuit architecture outside the design's context or independently of circuit switching.

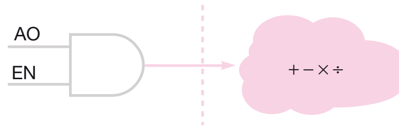

Figure 2 Operand isolation, data gating, or datapath gating can reduce power consumption, but it degrades timing.

To achieve the best result, you must reevaluate the architecture using a logic-synthesis tool, such as Design Compiler, employing the timing model and switching profile.

For example, consider a two-input multiplier with uneven switching activities on the operands. Although the multiplication is a commutative operation, the dynamic-power consequence is not. If you use the high-activity input for partial-product generation, the multiplier will consume more dynamic power due to a higher level of switching activities that propagate through the rest of the multiplier. If you switch the high-activity input to the input of the partial-product selector, you can lower the switching activity in the partial-product generator as well as the overall multiplier (**Figure 3**). This kind of optimization is hard to plan in the RTL-coding stage and is more suitable to perform during synthesis.

Applying this concept on a larger scale enables you to achieve more power savings. In general, irregularity in the data or a circuit provides a power-saving opportunity. For example, multimedia data usually has uneven activities among the data bits. It usually has lower activities in MSBs (most-significant bits) and higher activities in LSBs. If you are aware of this phenomenon, you can design datapath architectures so that the LSBs feed into the datapath tree downstream, hence reducing the dynamic power for audio- or video-signal processing. Likewise, you can use the circuit's irregularity to lower internal power and leakage power. For example, you can substitute regular cells with slower high-threshold voltage or low-drive cells whenever there is timing slack.

You can configure the DesignWare minPower architecture to create more timing slack to maximize this effect. However, manually exploiting the circuit's irregularity is difficult because it is imperative to balance the power cost against the area cost to avoid any adverse effect from over-aggressive power optimization. You must automatically consider timing and design needs during the architecture selection to realize the power savings with minimal area overhead.

The biggest advantage of power-friendly architectures is that they do not disrupt design flows. The power savings come from making power-smart choices when implementing the microarchitectures for the RTL code. This approach requires no changes to the higher-level software, system, or RTL design. After you select the architec-

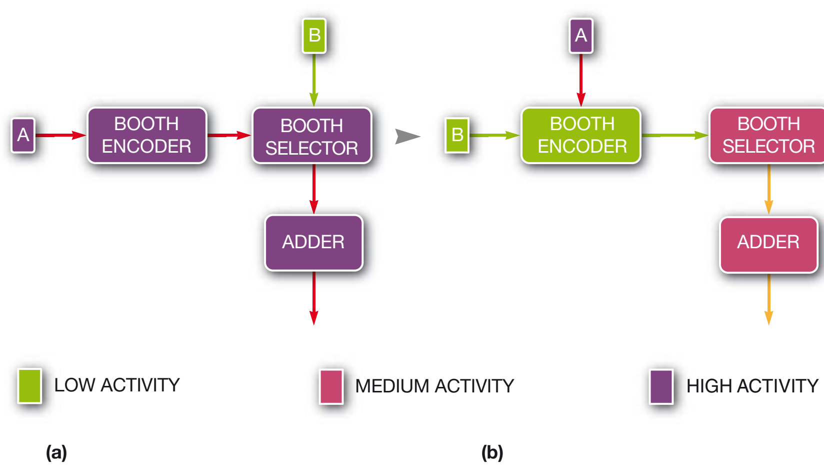

Figure 3 A high-activity input for partial-product generation causes the multiplier to consume more dynamic power (a). Switching the high-activity input to the input of the partial-product-selection block can lower the switching activity in the partial-product generator and the overall multiplier.

tures, the netlists go through the same gate-, physical-, or process-level optimization in the back end. You need not change design flows or design-database formats except for adding a new knowledge base—a synthetic-library database (.sldb file)—to the RTL-synthesis stage. The power savings increase the design project's original power strategy by as much as 42% additional power reduction at the block level and as much as 24% reduction at the chip level.

This architecture-level power-optimization approach does have some limitations, however. To get the power benefit, you must integrate in-house or third-party IP (intellectual property) into the design at the RTL because the optimization takes place at the RTL.

TO LOWER THE ENERGY CONSUMPTION OF YOUR NEXT SOC PROJECT, YOU MUST IDENTIFY WHICH PORTIONS OF THE SOC ARE CONSUMING THE MOST ENERGY.

The automatic IP insertion relies on a logic-synthesis tool, such as Design Compiler, to extract the datapath architecture from the RTL; therefore, the code must be in a style that the synthesis tool recognizes. In other words, if the datapath is in low-level RTL that already prescribes the architectures, the synthesis tool cannot alter the design's intent.

To enable architectural-level power optimization, designers should start from high-level RTL code using as much operator inference as possible. To allow extraction of larger datapath blocks, you should consider using automatic retiming instead of manually inserting a pipeline. Whenever possible, use a realistic representative switching profile, which usually improves the result, especially for applications that have unevenly distributed activities on the input.

Because power is a physical-domain characteristic, your standard-cell library can affect the power-optimization result. A standard-cell library with a collection of datapath cells that have good drive strength and threshold-voltage variations allows wider architecture selections.

Some libraries support special datapath cells but have few or no drive-strength variations or have them only with standard threshold-voltage implementations. You often do not select these cells, therefore limiting the number of available architectures. To improve results, use a standard-cell library with more drive strength and threshold-voltage variations that have accurately characterized power numbers.

You can't optimize what you can't observe. To lower the energy consumption of your next SOC project, you must first identify which portions of the SOC are consuming the most energy. It is worth distinguishing power consumption from energy consumption. To get a more energy-efficient design, you must pay attention to the circuits that remain on for a long time. Therefore, you must carefully analyze the power modes to identify the best energy-saving opportunities.

When working on these modules, decide early in the chip-planning stage to run these circuits at low clock frequencies and low voltage. The additional power-saving opportunities must come from designing more power-friendly circuits that require less switching activities and are built with a higher percentage of low-leakage and low-drive cells. The most inexpensive way of achieving this goal is by using power-friendly architectures that require no costly design-flow re-engineering efforts. EDN

AUTHOR'S BIOGRAPHY

Jay Chiang is product-marketing director at Synopsys, where he has worked for 11 years. He is responsible for managing the DesignWare library and the datapath-generator, JPEG, microcontroller, mobile-industry-processor interface, and memory-IP product lines. Before joining Synopsys, Chiang was a senior ASIC designer and chip architect at Xinex and a hardware-design engineer at Dynapro. He has a master's degree in management for science and technology from Oregon Health and Science University (Portland, OR) and a bachelor's degree from National Tsing Hua University (Hsinchu, Taiwan).

VENDORS TARGET OSCILLOSCOPE SWEET SPOT

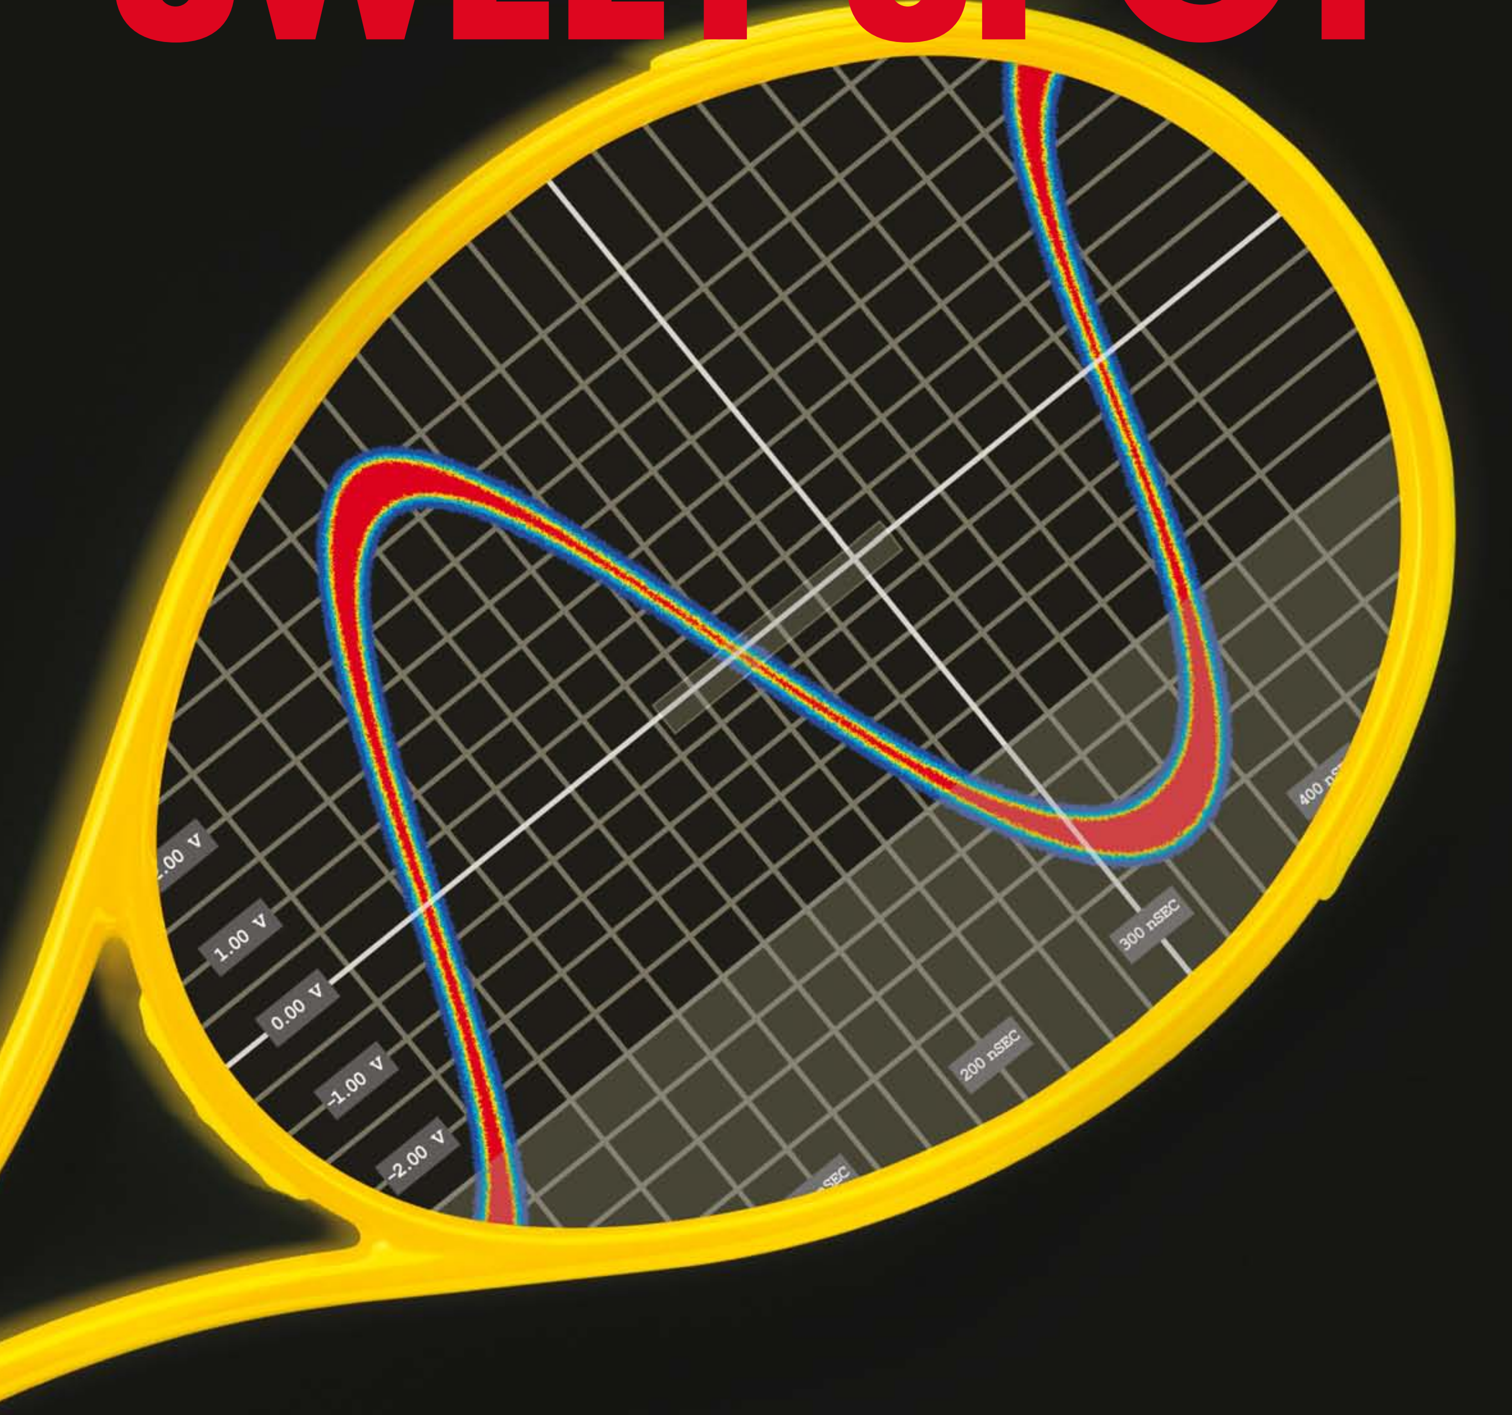

WHEN CHOOSING OSCILLOSCOPES IN THE 1- TO 4-GHz RANGE, ENGINEERS HAVE AN EXPANDING VARIETY OF PRICE, PERFORMANCE, AND USABILITY OPTIONS AS THE MARKET ACQUIRES A NEW COMPETITOR.

BY RICK NELSON • EDITOR-IN-CHIEF

Oscilloscopes having record-setting bandwidths garner trade-press attention as the major competitors leapfrog past each other (**Reference 1**), but instruments having maximum bandwidths of 1 to 4 GHz can serve many demanding applications. Vendors offering scopes with bandwidths in this range are also offering a variety of feature combinations, including triggering, waveform-capture capability, data-analysis capability, probing options, and user-interface functions, that help prospective customers find the price and performance levels they need for today's applications while providing head room for tomorrow's needs.

The market for 1- and 2-GHz oscilloscopes is so attractive that a new competitor, the Rohde & Schwarz test-and-measurement division, introduced models in that range in June. At a press conference announcing the instruments, Michael Vohrer, who was then chief executive officer and has since retired, said that the time-domain initiative represents an attempt of the privately held company to push into new markets and expand market share in its traditional markets for frequency-domain-analysis equipment. Vohrer placed the scope market at $1 billion and added that, with a highly diversified customer base, it represents lower volatility than do other segments.

MARKET-SHARE NICHE

Roland Steffen, head of the R&S test-and-measurement division, says that initial models in the new R&S RTO-oscilloscope line offer top bandwidths of 1 and 2 GHz. The company is not ignoring lower bandwidths, however, and introduced complementary RTM models with 500-MHz bandwidths. The RTO and RTM models combine to serve the 500-MHz to 2-GHz bandwidth range that is enjoying the largest share of market volume.

Prathima Bommakanti, senior research analyst for test and measurement at Frost & Sullivan, concurs with the perceived importance of that market niche. "Big giants, including Tektronix and Agilent, view the 500-MHz to 2-GHz range as a 'definite-demand' market," she says, referring to the bandwidth ranges that the new R&S scopes serve. Bommakanti's research indicates that there is constant demand for 500-MHz to 2-GHz scopes selling for $8000 to $20,000.

The new R&S instruments include RTO models in two- and four-channel versions with sampling rates of 10G samples/sec (**Figure 1**). The instruments support a Windows-driven touchscreen user interface. The 500-MHz RTM models offer 5G-sample/sec sampling and forgo the touchscreen interface but boot within 7 seconds to help provide fast measurement results. Prices for RTM instruments start at €5000, and prices for RTO instruments start at €12,000.

The new oscilloscopes don't represent Rohde & Schwarz's first corporate excursion into the time domain. Five years ago, the company acquired low-end scope-maker Hameg. Hameg will continue to supply instruments costing roughly €4000 and less through distributors, and the Rohde & Schwarz test-and-measurement division will serve the market for scopes operating at 500 MHz and more and selling for €4000 or more through its direct sales force.

Josef Wolf, head of the spectrum and network analyzers, EMC (electromagnetic-compatibility) tests, and oscilloscopes subdivision at R&S, commented during the June press conference on the development effort that went into the new scopes. That effort focused on the high-level integration of analog, mixed-signal, and digital subsystems. A key goal was a low-noise analog front end, which the company achieved through the use of a single-core SiGe (silicon-germanium), 10-GHz ADC with an ENOB (effective number of bits) better than seven. A 90-nm ASIC with 15 million gates provides hardware implementation of digital-signal-processing functions, enabling the analysis of 1 million waveforms/sec.

The 2-GHz top-of-the-line RTO models employ a purely digital trigger system that eliminates the alignment errors that can occur with software-compensation schemes with separate analog triggers. The company specifies the RTO models' trigger jitter in femtoseconds rather than picoseconds. In addition, the digital trigger eliminates rearm times associated with analog triggers, which can mask events of interest that occur shortly after an analog trigger. The RTO provides as much as 20 times less blind time than competitive models to help identify intermittent problems, according to Wolf.

The market will decide how much share the new R&S scopes will gain with their price, performance, and features, and Bommakanti at Frost says that a clearer picture will emerge in

2011. Nevertheless, she expects the instruments to be competitive. "Market participants opined that they are seeing Rohde & Schwarz already competing vigorously in this market," she explains. "From our perspective, having researched the general-purpose-test-equipment market for a number of years, we believe R&S can leverage its brand reputation based on quality in the overall test-and-measurement-equipment market. Oscilloscopes were just the missing pieces of their product line." According to Steffen at R&S, the company's first efforts in oscilloscope marketing will be to target its customers for frequency-domain equipment who also need oscilloscopes in these bandwidths; the company will then look beyond its customer base.

AT A GLANCE

- According to an analyst at Frost & Sullivan, oscilloscope vendors see a consistent demand for 2-GHz oscilloscopes.
- Bandwidth is only one factor; noise performance, memory depth, and data-analysis capabilities are also important features.
- Scopes need to connect to target systems with an arsenal of probes.
- Oscilloscope usability is a subjective factor that's difficult—but not impossible—to quantify.
- The market for higher-performance scopes will experience higher growth rates, but lower-bandwidth models will dominate in unit shipments.

SERIAL-DATA OPTIONS

Meanwhile, the company's competitors are not standing still. They have their own ideas of the price, performance, and feature combinations that they believe make their scopes with comparable bandwidths competitive. "In terms of performance, oscilloscopes can no longer be quantified by bandwidth alone," says Joseph Ting, product manager for high-frequency instruments at Yokogawa Corp of America. "Many other common measures of oscilloscope performance include noise/accuracy, frequency-response curve, waveform-acquisition rate, memory depth, memory handling, real-time- and postacquisition-analysis capabilities, and responsiveness under load." Yokogawa offers scopes having bandwidths as high as 1.5 GHz, including the DLM6000 series digital- and mixed-signal oscilloscopes. They can perform waveform characterization, include tools for detecting glitches and anomalies, incorporate signal-enhancement and noise-reduction technologies, and come with a range of options for serial-bus analysis and power measurement. Models offer four channels plus 16- or 32-bit logic inputs and feature 500-MHz, 1-GHz, or 1.5-GHz bandwidths.

The variety of potential options that 1- to 4-GHz scopes are able to incorporate gives vendors room to maneuver as they compete, and they continue to introduce features. For example, LeCroy last month introduced a series of serial-data-product enhancements for all Windows-based WaveSurfer, WaveRunner, WavePro, and WaveMaster scopes. The serial-data enhancements include an ARINC (Aeronautical Radio Inc) 429 D decoder option, which, coupled with the company's MIL-STD (military-standard)-1553 TD (trigger-decode) package, rounds out its military- and commercial-aviation enhancements, according to Bill Driver, product-marketing manager at LeCroy. LeCroy is also introducing support for MIPI (Mobile Industry Processor Interface) and the more than 10 standards it encompasses, including DigRF. Last month's introductions include a package of serial-protocol measurement, data-extraction, and graphing tools that apply a specialized set of timing and graph parameters to standards such as I^2C (inter-integrated circuit), SPI (serial-peripheral interface), UART (universal asynchronous receiver/transmitter), RS-232, CAN (controller-area network), LIN (local-interconnect network), FlexRay, and AudioBus (**Figure 2**).

Figure 1 Rohde & Schwarz has jumped into the oscilloscope market with its RTO models, which come in two- and four-channel versions with bandwidths of 1 and 2 GHz. The units' sampling rate is 10G samples/sec. The instruments support a Windows-driven touchscreen user interface.

SCOPES NEED ADAPTABILITY

The many serial-bus options available for scopes such as LeCroy's WaveRunner suggest that the instruments must be adaptable for a variety of applications, and serial-bus options aren't the only areas in which the scopes exhibit flexibility. "An oscilloscope in the 1- to 4-GHz class would still be qualified as a general-purpose testing tool that should have the capability to adapt to a wide range of applications," says Driver. LeCroy's scopes in this range include an integrated switchable 50Ω and 1-MΩ front-end termination. The 1-MΩ option allows you to use simple passive probes for applications that require a basic understanding of a signal without concern for noise or timing accuracy. In contrast, the 50Ω path can serve in applications that require the most accurate signal shape, including cabled systems and those employing single-ended FET probes or active differential probes.

Joel Woodward, senior product manager for Agilent's oscilloscope group, agrees that flexibility is an important consideration in oscilloscopes that span the 1- to 4-GHz range. "Engineers in this segment typically encounter a wide number of issues that they need a scope to help solve," he says. Agilent serves this segment with its Infiniium 9000 series family (**Figure 3**), which offers bandwidths of 600 MHz to 4 GHz, enabling Agilent scope customers to pick the price and performance that match

their needs. The scopes come with a selection of debugging, protocol-triggering and -decoding, jitter-analysis, and compliance applications. Agilent offers protocol-triggering and -decoding support for 12 protocols, including I^2C, SPI, RS-232/UART, CAN, LIN, FlexRay, JTAG (Joint Test Action Group), USB (Universal Serial Bus), PCIe (Peripheral Component Interconnect Express), MIPI D-PHY (digital physical layer), SATA (serial advanced-technology attachment), and 8b/10b.

According to Woodward, Agilent is the only vendor to support both traditional hardware triggering and software-based triggering, which the company calls InfiniiScan. Hardware-based triggering enables a wide range of predefined trigger conditions from which users can select. The hardware triggering embraces analog, digital, and protocol triggers to catch even the rarest of events. "Software-based triggering enables the user to graphically define a trigger," Woodward says. "The scope then compares each acquisition to the trigger specification and displays only the acquisitions that match the trigger specification." You can also cascade hardware and software triggers to create multistage triggers.

Digital or logic triggering is helpful in systems for which you want to verify performance across many data points but have few concerns about signal integrity, according to Chris Loberg, senior technical-marketing manager at Tektronix. "The digital trigger can speed up verification without tying up time worrying about analog-signal characteristics that may not be important due to a lower clock speed or logic type," he says. Analog event-based triggering is critical, however, for performing signal-integrity checks or for drilling deeper into signal performance during debugging. When you combine analog triggering with a high waveform-display rate, such as the one that Tektronix's DPX (digital-phosphor technology) provides, the debugging capability improves dramatically, Loberg adds.

Figure 2 Windows-based oscilloscopes, such as LeCroy's 2-GHz WaveRunner (left), can accommodate serial-data packages that support ARINC 429, MIPI, I^2C, SPI, UART, RS-232, CAN, LIN, FlexRay, and AudioBus interfaces. LeCroy's new measurement-and-graphing package (right) displays the left- and right-channel data from a digital-audio source that has undergone conversion into corresponding analog values and plotting over time to represent the converted analog-audio waveform. The example is from a Pink Floyd audio track and represents about a half-second of audio.

Woodward echoes Driver's emphasis on the importance of flexible probing options. "Scopes need to connect to target systems," he says, and customers must choose the right probes and be able to later add probes. Agilent's Infiniium scopes work with a range of single-ended, differential, high-voltage, and current probes. "We've developed this probing arsenal over a number of years by working with key customers and helping them solve their probing challenges," he explains. "During just the last year, Agilent introduced more than 27 probes."

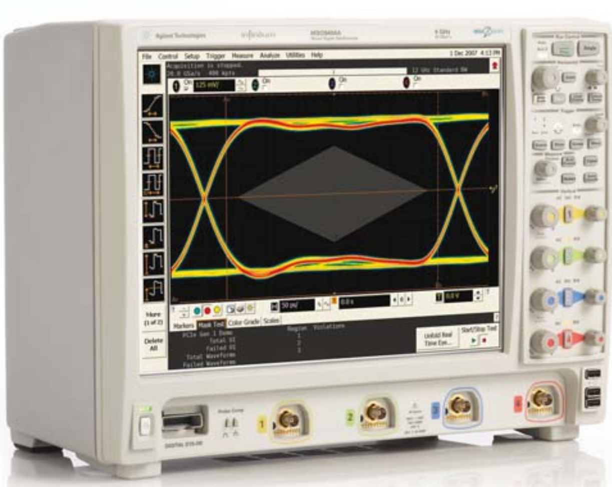

Figure 3 Agilent's Infiniium 9000 series scopes offer bandwidths of 600 MHz to 4 GHz. Each model comes in a 9-in.-deep package and has a 15-in. LCD.

USABILITY IS KEY

Tektronix addresses the 1- to 4-GHz scope market with scopes such as its DPO7000 series (**Figure 4**) and its MSO (mixed-signal-oscilloscope) 4000 series. "One thing that commonly happens is users tend to focus in on one banner spec or feature rather than the full breadth of factors, including value, performance, usability, probing solutions, and vendor expertise," says Loberg. Usability is a key factor. "When deciding to invest in this level of scope, the usability of the instrument is critical. For example, many of today's modern scopes in this range feature serial decoding that quickly speeds up an engineer's ability to understand the exact traffic being sent on common serial buses," he adds.

Ting at Yokogawa agrees. "User interface is very important not only for engineers' productivity, but also for their mental health. An intuitive interface, menu tree, and optimized keystrokes for each operation all make for an improved user experience." He notes that usability often depends on familiarity. Each vendor implements user-interface methods unique to its instruments. Yokogawa oscilloscopes, for example, all employ a jog shuttle, which comprises an inner dial with detents and an outer dial with a spring-loaded ring. "Users who are familiar with our interface are reluctant to change to other vendors [and vice versa]," Ting says.

Woodward at Agilent attributes the usability features of its Infiniium scopes to the company's continual refinement of the instruments, which it bases on input from a large installed base. In response to Agilent's queries about us-

ability, customers respond, "Please don't change it," according to Woodward. Agilent scopes ranging from the Infiniium models to the real-time-bandwidth-champion, 32-GHz model employ the same software architecture.

LeCroy's Driver also touts usability. "We have pioneered a lot of user-interface techniques," he says, including a touchscreen, multiple grids for independently viewing waveforms, and the ability to draw a box on the screen to zoom in on a waveform. "When a user interface does what customers expect it to do, they spend more time debugging instead of climbing the learning curve," he adds. Like Woodward, Driver cites a consistent architecture across product lines: "When a customer picks up a 400-MHz LeCroy oscilloscope or a 30-GHz LeCroy oscilloscope, the user interface is the same."

Usability, however, can be a tough sell. "The usability is equally important [as other features] but unfortunately subjective," says Driver. The decision to purchase a scope in this class typically begins with a banner-specification comparison, then a price comparison, and a usability evaluation, he explains. Ting notes that usability can be a key deciding factor, but it would be difficult to quantify it to an actual cost value.

Tektronix's Loberg also emphasizes usability and an effective user interface in today's fast-paced lab environment. "Engineers need familiar, intuitive instruments that save them from having to spend time adapting to the instrument's operation," he says. "Many are jumping from a workstation-based environment to the lab and want to solve problems or verify performance quickly."

Although you cannot reduce usability to a banner spec, such as bandwidth or number of channels, Tektronix has made an effort to quantify usability, at least as it relates to performing a set of tasks, according to Loberg. The company commissioned Hansa Research to conduct a study asking users, given a test circuit, to set up a Tektronix MSO4000 series oscilloscope and competing oscilloscopes to monitor for glitches and runts, set up a trigger and capture a runt, and search the waveform to locate all instances of the runt (**Reference 2**). The study found 53% improvement in debugging time with the use of the Tektronix scope.

The banner specs themselves are also open to misinterpretation. "A common mistake is to specify a bandwidth that is not sufficient to see all of the signal content," explains Driver, who notes that customers often ask for bandwidth equal to the bit rate of the signal under test. "A common rule for specifying bandwidth is to have a bandwidth equal to or greater than five times the primary frequency or … the fifth harmonic," he says.

Ting echoes that advice, saying that an oscilloscope that can acquire signals up to the fifth-order harmonics can generally reconstruct a pulse signal with sufficient accuracy. He also cautions that some users may overlook the fact that the oscilloscope's sample-rate specification is only a maximum. Both memory depth and observation time restrict the actual sample rate.

SPECS, FUTURE MARKETS

Driver advises engineers to consider future needs when evaluating banner specs. If the current budget situation prevents the purchase of a higher-bandwidth oscilloscope, he suggests the purchase of an upgradable product. Woodward at Agilent has similar advice, noting that Agilent customers can buy a 1-GHz model now and upgrade to 2.5 or 4 GHz as their needs grow. Similarly, they can easily upgrade to add 16 digital channels, he says.

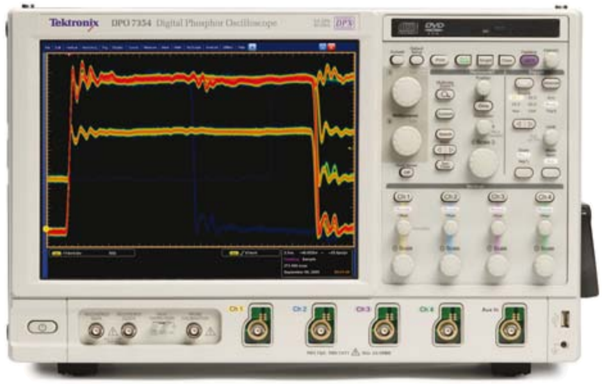

Figure 4 Tektronix's DPO7000 series oscilloscopes support the Pinpoint triggering system, which provides a suite of advanced trigger types on both the A and the B triggers. Pinpoint triggering offers more than 1400 triggering combinations.

Once you determine specs such as the bandwidth you need and can afford, it's relatively easy to compare them on a data sheet. Other specs, although quantifiable, are open to misinterpretation. "Users should be aware that waveform capture/update rate is directly related to their ability to detect intermittent anomalies that cannot be perfectly isolated by triggers," says Yokogawa's Ting.

Woodward cites an example regarding waveform-capture rate, which can vary depending on scope settings. "The Infiniium 9000 delivers 250,000 waveforms per second in segmented-memory mode and up to 3000 waveforms per second with shallow memory," he says. "With 1M point of memory turned on, the scope still delivers more than 600 waveforms per second."

Frost's Bommakanti predicts steady, low-single-digit growth for scopes in the 2-GHz range. Revenue growth should be slower for midrange scopes than that for high-performance scopes, but unit shipments of lower-bandwidth models will continue to experience higher unit shipments. Expect vendors to continue to fine-tune performance to capture their fair share of the market.EDN

FOR MORE INFORMATION

Agilent Technologies
www.agilent.com

Frost & Sullivan
www.frost.com

Hansa GCR
www.hansagcr.com

LeCroy
www.lecroy.com

Rohde & Schwarz
www.rohde-schwarz.com

Tektronix
www.tektronix.com

Yokogawa Corp of America
tmi.yokogawa.com

REFERENCES

1 Rowe, Martin, "Agilent introduces 32-GHz oscilloscope," *Test & Measurement World*, April 27, 2010, http://bit.ly/dyfAyb.

2 Nelson, Rick, "Testing embedded data buses and analog signals," *EDN*, June 24, 2010, pg 36, http://bit.ly/aRFBze.

You can reach Editor-in-Chief Rick Nelson at 1-781-869-7970 and richard.nelson@cancom.com.

DESIGN NOTES

A Low Power, Direct-to-Digital IF Receiver with Variable Gain

Design Note 482

Walter Strifler

Introduction

Modern communication receivers require an ADC to digitize an incoming analog signal for decoding in a suitable FPGA device. The direct-conversion method of receiver design typically performs a single frequency downconversion and an analog-to-digital conversion (ADC) near baseband. While elegant and simple, this receiver architecture has problems with in-band blockers, out-of-band interferers and LO leakage reflections within the receiver itself.

In the face of these problems, basestation receivers often require a robust solution that is achieved using tried-and-true system methods of downconversion to an intermediate frequency (IF) in the range of 70MHz to 240MHz. Demodulating and decoding the IF signal can be performed by various means, but an increasingly popular and cost-effective method is direct-to-digital IF conversion using the recent generation of high speed, low power pipeline data converters available from Linear Technology.

This design note describes a variable gain amplifier plus analog-to-digital converter (VGA + ADC) combination circuit that preserves the IF receiver dynamic range over a 31dB gain adjust range and effectively demodulates and digitizes both the I and Q information in a single step. The combination LTC®6412 VGA and LTC2261 14-bit ADC circuit subsamples a 140MHz WCDMA IF channel at 125Msps and provides an equivalent input NF and IP3 that rivals some of the best laboratory spectrum analyzers, while consuming less than 0.5W of power.

IF Receiver Performance

The performance of a demonstration receiver circuit is shown in Figure 1. The inset graph of Figure 1 shows the noise-like distribution of the WCDMA signal and is similar to CCDFs of other modern communication signals. At VGA maximum gain, the signal generator power is adjusted to $-12.5dBFS_{RMS}$ to occupy most of the ADC code range without clipping.

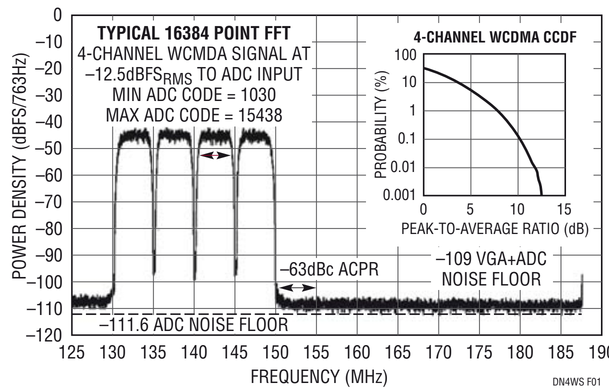

Figure 1. Typical WCDMA Performance at $-12.5dBFS_{RMS}$ over All Gain Adjust Settings. Insert Shows WCDMA CCDF

As the input signal power is adjusted higher, the VGA gain is adjusted down to maintain $-12.5dBFS_{RMS}$ and simulate the automatic gain control (AGC) response of a typical receiver. The FFT of the digitized receive signal is plotted over a full Nyquist zone and exhibits a 63dBc ACPR with no measurable spurs and only 2.6dB degradation in the ADC noise floor, *over the full 31dB gain adjust range.* This represents an effective input NF of 13dB and input IP3 of 23dBm for the VGA + ADC pair at 140MHz at maximum gain. The (IP3-NF) delta of 10dBm determines the effective dynamic range of the receive pair and is nearly constant over the entire gain adjust range.

Measurement Details and Receiver Circuit

An Agilent E4436B source generates the multichannel WCDMA test signal with a typical adjacent channel power ratio (ACPR) of 50dBc to 55dBc, perfectly adequate to meet the WCDMA system specifications but insufficient to demonstrate the full quality of this VGA + ADC combination. The test signal is amplified with a high linearity Triquint AH202 and sharply filtered with a SAWTEK 854920 to reduce the test signal's ACPR skirts below 65dBc.

The WCDMA signal is representative of the wideband, noise-like signals found in modern communication systems such as LTE, 802.11g, and WiMAX to name a few. Interestingly, this convergence of statistical signal behavior was predicted over 60 years ago in Claude Shannon's communication theory. He found that the methods to increase spectral efficiency in a modulation format will, by necessity, exercise many degrees of freedom in signal space and approximate the process of additive white Gaussian noise. This was an amazing insight considering the simple AM and FM signals of Shannon's day. This is also a practical insight. One representative noise-like signal can be used to characterize an RF receiver and estimate the performance of other noise-like signals.

Figure 2 details a receiver circuit optimized for a 140MHz center frequency and a 20MHz bandwidth typical of a 4-channel WCDMA signal. The filtered test signal feeds to the VGA input balun to perform a single-ended to differential conversion at the input of the LTC6412. The LTC6412 output connects to a simple tank circuit and RC network at the input of the LTC2261. This matching circuit routes bias current to the VGA while performing a low Q impedance transformation to the 100Ω differential load. The matching circuit and RC load also serve to dissipate the differential and common mode charge injections emanating from the sampling switches at the ADC input. This is an important consideration, as these charge impulses need to dampen to better than –85dB during a sampling window (4ns) to preserve the full spur-free dynamic range (SFDR) of the LTC2261. The better damping circuits tend to be small and tight to avoid unnecessary reflection delays and mismatch between the VGA output and ADC input. This particular matching circuit uses 0402 components for most elements and fits inside a board area of 5mm × 10mm.

The balance of connections to the VGA and ADC follow the recommendations of their respective data sheets. The LTC2261 14-bit ADC runs off 1.8V and consumes 127mW at 125Msps. The LTC6412 VGA runs off 3.3V and consumes 360mW for a total power consumption of 490mW.

Conclusion

The LTC6412 VGA drives the LTC2261 14-bit ADC with little compromise in the ADC performance. The VGA buffers the ADC sampling input and provides 31dB of gain adjust to expand the effective dynamic range of the subsampling IF receiver. The LTC2261 is part of a family of 12- and 14-bit low power data converters designed for maximum sampling rates in the range of 80Msps to 125Msps. For complete schematics of this receiver, visit the LTC6412 or LTC2261 product pages at www.linear.com.

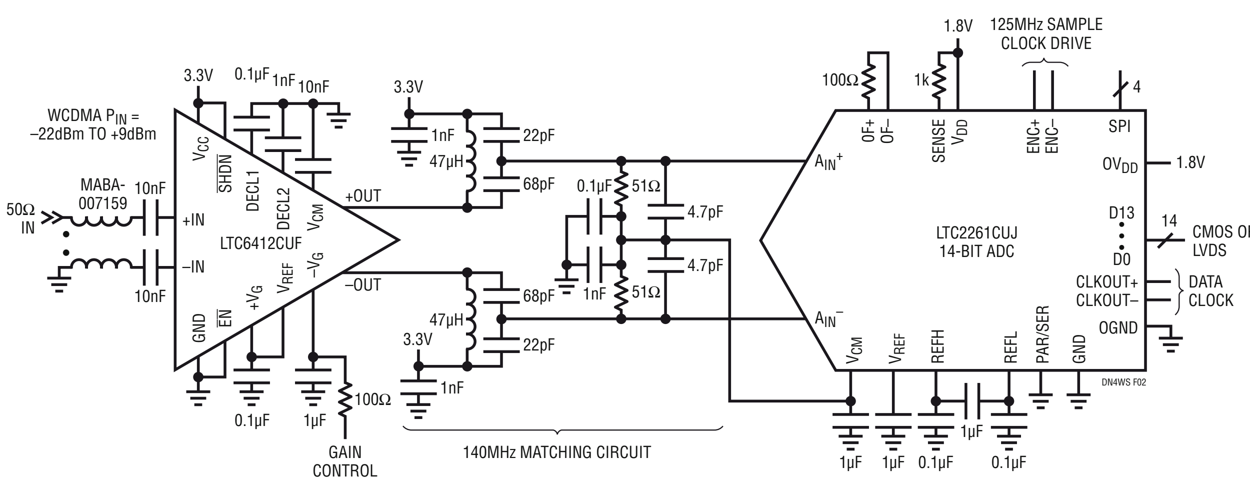

Figure 2. VGA + ADC IF Receiver Circuit. Supply Decoupling Capacitors to the VGA and ADC Omitted for Clarity. For This Measurement, the LVDS Bus Connects to Linear's Data Acquisition Board DC890B for Computer Control and Data Analysis

Data Sheet Download

www.linear.com

For applications help,
call (408) 432-1900, Ext. 2525

Linear Technology Corporation
1630 McCarthy Blvd., Milpitas, CA 95035-7417
(408) 432-1900 • FAX: (408) 434-0507 • www.linear.com

dn482f LT/TP 0910 226K • PRINTED IN THE USA

BY KRIPASAGAR VENKAT • TEXAS INSTRUMENTS

Glass-breakage detector uses one microcontroller

A GLASS-BREAKAGE DETECTOR CAN DETECT WHEN A WINDOW OR DOOR BREAKS IN A HOME OR BUSINESS, SERVING AS A MONITORING DEVICE TO ENHANCE SECURITY BY DETECTING ILLEGAL ENTRY.

A glass-breakage detector works either independently or in conjunction with other antitheft devices to form a security system. The detector essentially captures and analyzes any acoustic activity and reports whether glass breakage has occurred. Due to their mode of operation, these detectors depend heavily on the quality of sound events, posing numerous challenges to the designer. The detector must also be able to reject all failure alerts—sounds that are not true glass breakage. This article discusses an efficient and robust glass-breakage detector using a low-cost microcontroller.

Microcontrollers are low-end processors that find use in applications such as simple digital real-time clocks and complex smart-metering systems. Microcontrollers suit these applications because they cost less, consume less power, and are easier to use than most other types of digital processors. In simple applications with limited requirements, it is easy to achieve low cost and low power. However, with the trend toward using microcontrollers in complex applications, it becomes a challenge to maintain low cost and achieve low power. Engineers must now try to get the best performance with the lowest possible cost. To achieve this goal, they face microcontroller-architecture restrictions, such as lower on-chip memory, a limited peripheral set, lower operational speed, and a smaller pin count. Engineers must optimize everything these microcontrollers offer for use in fairly complex applications, such as the glass-breakage detector.

DESIGN CONSIDERATIONS

A robust glass-breakage-detection algorithm should be able to easily distinguish valid glass breakage from other sound events. All glass-breakage-detection algorithms capture sound events, analyze their time and frequency components, and make a decision. Glass-breakage sounds vary by type of glass, thickness, acoustic environment, distance, the object the would-be thief uses to break the glass, and other factors. All glass-breakage-detection algorithms are inherently similar but vary slightly depending on conditions, so one algorithm will not work for all conditions. Installers usually fine-tune the algorithm during final installation in a home or business.

You can analyze a valid glass-breakage signal in the time domain or the frequency domain. **Figures 1** and **2** show a typical glass-breakage signal in the time and the frequency domains, respectively. This sound falls well within the audio spectrum of 20 Hz to 20 kHz. The time-domain waveform relates to the sound that a listener hears, and the frequency-domain waveform gives the complete frequency content of the signal. These plots provide valuable information in the design of an efficient algorithm for breakage detection. The time-domain plot indicates that the waveform is dense and that a lot of activity occurs in short intervals. This activity relates to the fact that the signal contains a lot of high-frequency components and that the waveform has a lot of ze-

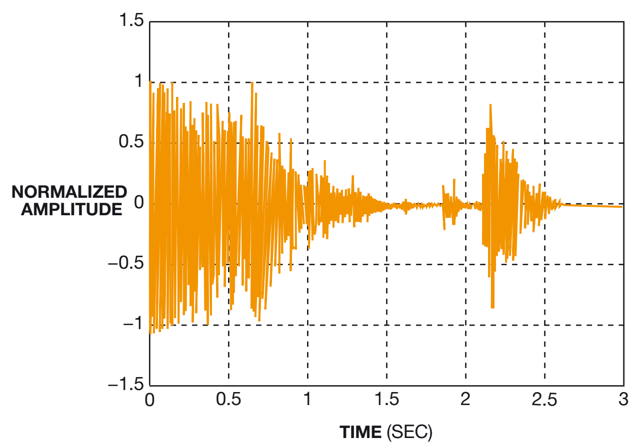

Figure 1 A typical glass-breakage signal in the time domain falls well within the audio spectrum of 20 Hz to 20 kHz and relates to the sound that a listener hears.

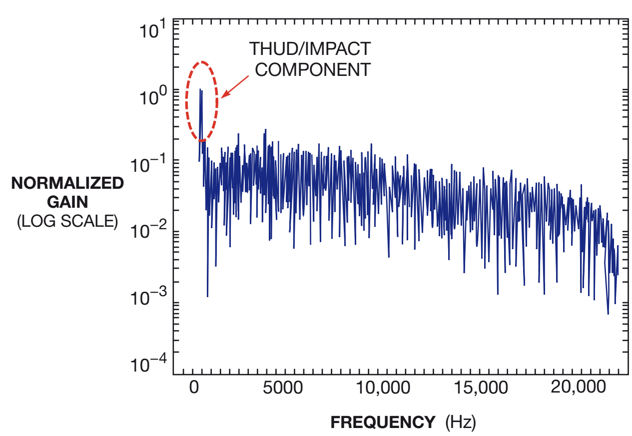

Figure 2 The frequency-domain waveform gives the complete frequency content of the signal.

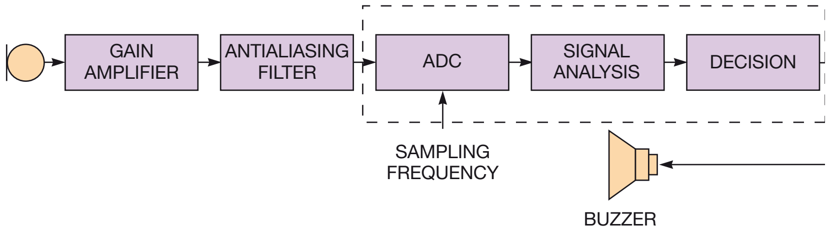

Figure 3 A microphone captures sound events, and an antialiasing filter following a gain amplifier handles signal amplification and filtering of high-frequency components.

ro crossings and high peaks. These characteristics, although they provide good information, seem to mimic white noise, and it is a challenge for designers to distinguish between the waveform's characteristics and white noise.

The frequency response involves similar challenges. The glass-breakage signal's components spread over the entire spectrum with fairly equal energy, which is typical of white noise. However, a peak occurs at approximately 200 to 300 Hz; this peak—the frequency component of the sound caused by the initial impact to the glass during breakage—provides the much-needed distinction. The impact is a low-frequency signal among all the high-frequency glass-breakage sounds that follow. You can view this impact, or thud, as the sound that occurs when an object hits the glass. It is difficult to recognize this information in the time-domain waveform, but you do know that this sound precedes all other sounds during breakage.

SYSTEM COMPONENTS

A glass-breakage detector must always be on and should be able to process any sound activity in real time. However, you can turn off some of the detector's blocks or put them into low-power modes when they are not operating. A microphone captures sound events, and an antialiasing filter following a gain amplifier handles signal amplification and filtering of high-frequency components (**Figure 3**). The antialiasing filter rejects any frequencies above the audible range of 20 kHz and avoids violating the Nyquist criterion during digitization of an analog signal.

The blocks within the dashed lines are parts of the processor, which can be an ASIC, a microcontroller, or a DSP. The ADC converts the analog signal to a digital sample for processing in the digital domain. The sampling frequency, F_S, depends on the frequency content of the signal. Because this circuit uses a 20-kHz antialiasing filter, the sampling rate must be at least 40 kHz to preserve the original signal's content and integrity. The signal-analysis block encompasses all of the signal processing necessary for the detection or rejection of glass breakage. Once this detection completes, the decision block activates an indicator, such as an LED or a buzzer, to indicate glass breakage.

HARDWARE SPECS

Most glass-breakage detectors operate from batteries; for sufficient battery life, the design should have low power consumption. The choice of all hardware components in this design depends on their ability to contribute to the low-power design. An analog signal starts at the microphone and ends at the ADC. The choice of the microphone is important because its performance contributes to the success of any glass-breakage detector. The microphone should also be able to capture and preserve key sound components, such as impact and other high-frequency components that the detector algorithm uses. The microphone must be on most of the time to capture any sound activity and hence must consume less power to bring down the overall system current. The gain amplifier—usually an operational amplifier in inverting or noninverting mode—has a gain higher than unity. The op amp must provide sufficient gain to the sound—on the order of tens of millivolts—the microphone captures. The op amp must always remain on and must have a small turn-on current. The antialiasing filter is also an op amp, filtering in the analog domain, and is usually a simple first- or second-order-unity-gain lowpass filter.

The most important choice in this design is that of the signal processor. You can use ASICs, microcontrollers, or DSPs, depending on the application. Like smoke detectors, most glass-breakage detectors are placed inside homes or offices at locations that ensure security and safety. However, they require battery power so that you can place them anywhere without worrying about their proximity to power outlets and to ensure that they will continue to operate in the absence of power on the mains.

You must choose a low-power, programmable, easy-to-use,

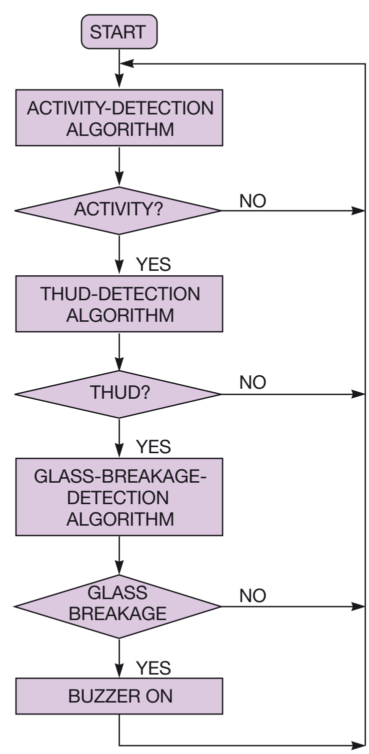

Figure 4 The high-level software flow of the algorithm includes activity detection, thud detection, and glass breakage in the order of occurrence in time.

and inexpensive processor with good processing capabilities for real-time operation. The microcontroller is the best of these choices because it meets all of these requirements. Some microcontrollers also integrate analog peripherals, which further reduce overall system cost.

TABLE 1 CURRENT AND TIMING CONSIDERATIONS

Condition/mode	Peripherals on	Clocks	Current	On time
Low-Power Mode 3	Timer on (up mode)	MCLK=DCO=off SMCLK=DCO=off ACLK=VLO≈12 kHz TACLK=ACLK=VLO	0.6 μA	2.5 msec
Activity detection/ AM1 (Active Mode 1)	Microphone on Op amp 0 on	MCLK=DCO=12 MHz SMCLK=DCO=12 MHz ACLK=VLO≈12 kHz	4.8 mA	20 μsec
Thud detection/ AM2 (Active Mode 2)	Microphone on Op amp 0 on Op amp 1 on Timer on (up mode) ADC 10 on	MCLK=DCO=8 MHz SMCLK=DCO=8 MHz ACLK=VLO≈12 kHz TBCLK=SMCLK=8 MHz	4 mA	32 msec
Glass-breakage detection/ AM3 (Active Mode 3)	Microphone on Op amp 0 on ADC 10 on	MCLK=DCO=12 MHz SMCLK=DCO=12 MHz ACLK=VLO≈12 kHz	5.8 mA	60 msec

SOFTWARE SPECS

An antialiasing filter with a 20-kHz cutoff frequency filters the analog signal from the microphone. To digitize this signal, the sampling rate must be greater than 40 kHz, and the ADC must be able to support that rate. For real-time operation, the filter must complete the required processing between successive-sampling instants. For example, if the maximum CPU frequency is 12 MHz, the number of CPU cycles between successive samples is only 300, which is tight for signal processing. You can choose a processor that supports a higher CPU clock for increased CPU cycles. Doing so would increase power consumption, however, and therefore decrease battery life. Hence, one must make a trade-off between the algorithm's complexity level and battery life.

The thud occurs at the beginning of a glass-breakage sound. This thud signal is present in most sounds, such as the noise that occurs when a door or cabinet closes, an object hits the ground, or someone claps his or her hands or knocks on a door. However, these sounds lack the high-frequency components of typical glass-breakage signals. Other sounds, such as the noise a coffee grinder makes, loud music, motorcycles in motion, or the noise a wineglass makes when it breaks, have similar high-frequency components but no thud component.

The glass-breaking algorithm exploits the fact that these two types of components are on either side of the frequency spectrum and occur independently of each other in time.

GLASS-BREAKAGE ALGORITHM

Figure 4 shows the high-level software-flow diagram of the algorithm, including activity detection, thud detection, and glass breakage in the order of occurrence in time. Approximately every 2.5 msec, the microphone and op amp 1 turn on to check for any sound activity. In the absence of any significant activity, they turn off, and the microcontroller goes into a low-power state. If significant activity occurs, the software proceeds to thud detection, during which the ADC turns on, followed by signal processing to check for the thud component. If a thud is present, the algorithm proceeds to glass-breakage detection. Otherwise, the algorithm reverts back to activity detection. If glass-breakage detection is successful, an onboard LED or buzzer activates to indicate this event. The glass-breakage detector then reverts back to activity detection.

Activity detection compares the ADC's input values to prefixed thresholds on either side of zero to distinguish

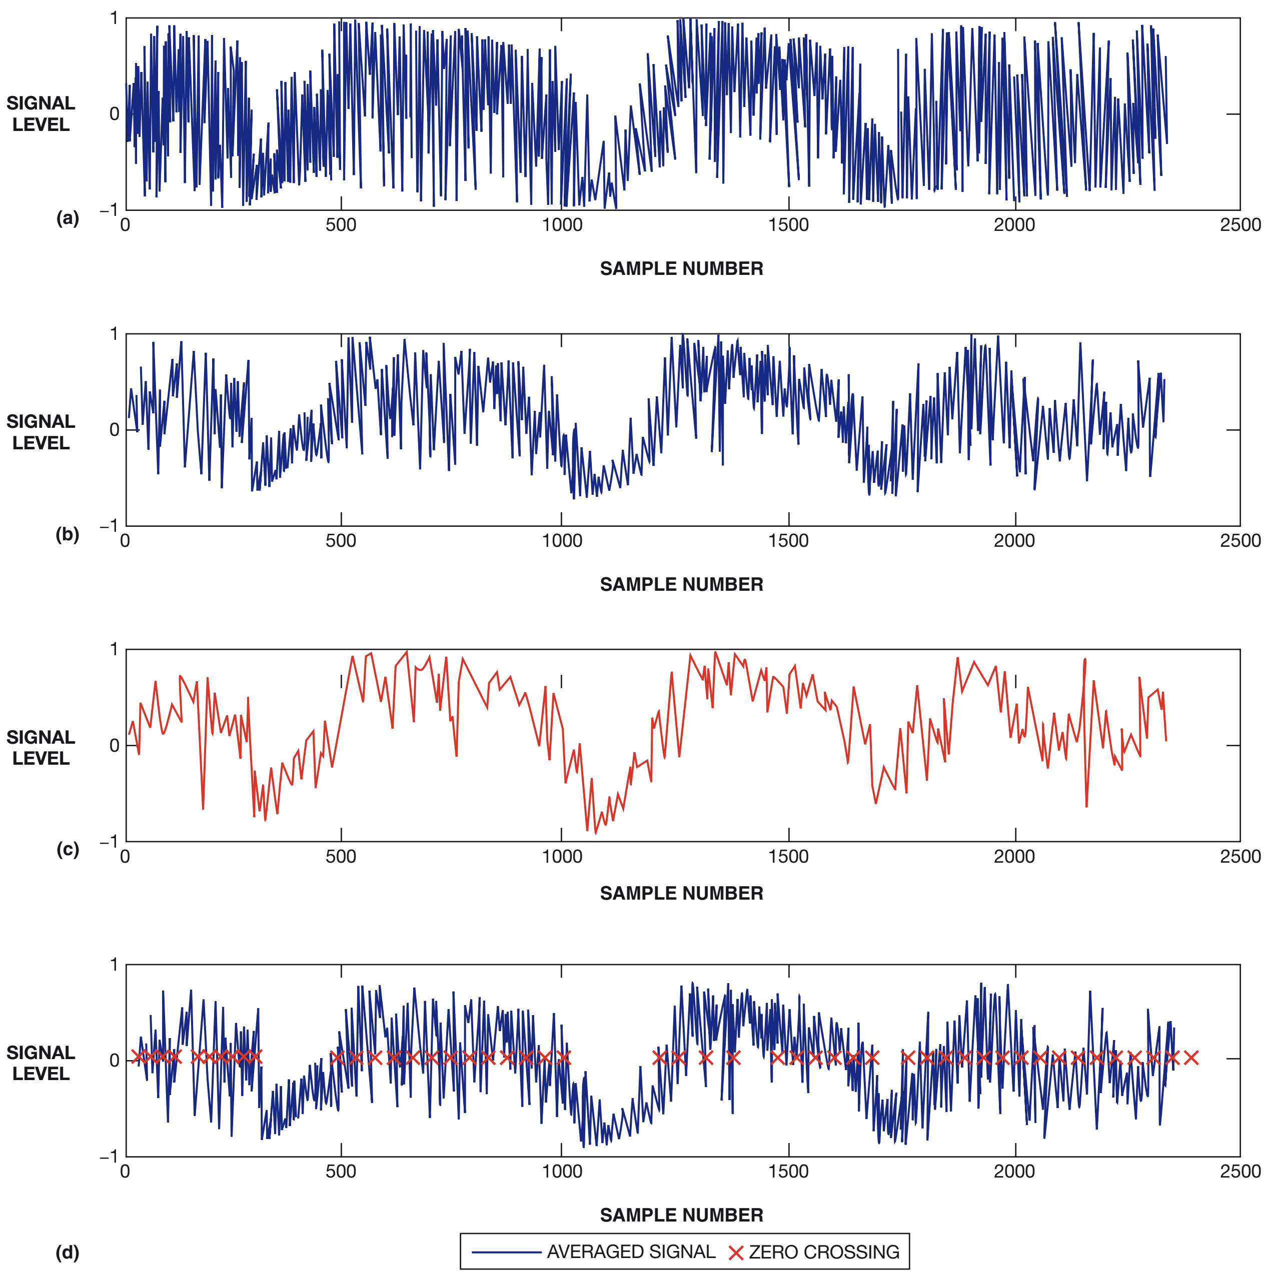

Figure 5 The first stage of processing occurs on every sample once the algorithm detects a thud. This stage uses a 20-kHz antialiasing filter and increases the ADC's sampling frequency to 40 kHz.

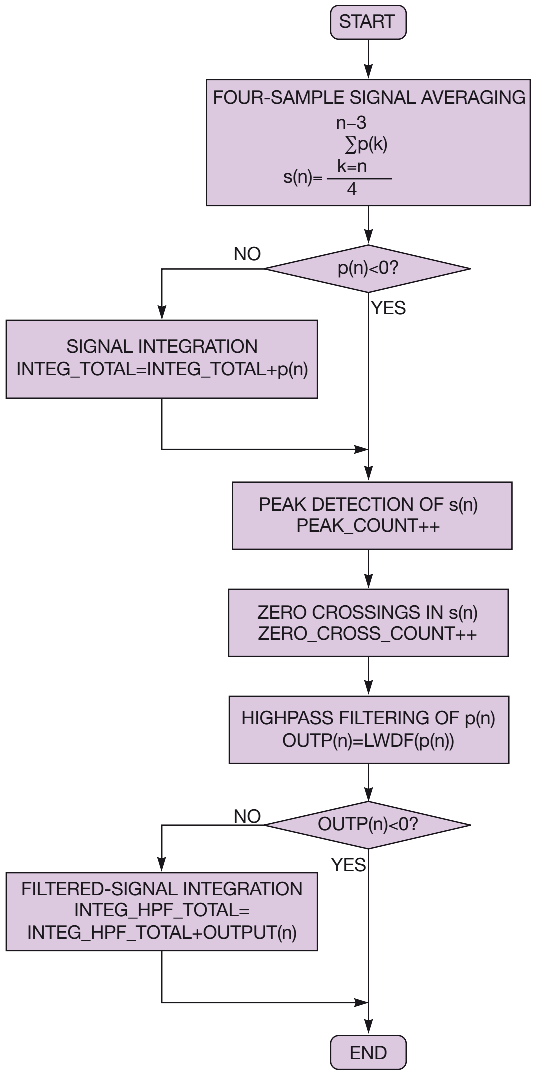

Figure 6 The operations for this stage include signal averaging, zero-crossing detection, and peak detection, which occur for approximately 60 msec, or approximately 2400 samples.

a true signal from noise. The thud component occurs only during initial impact, and a digital lowpass filter with a cutoff frequency of 350 Hz filters only the first few samples of the incoming signal. The system accumulates and averages the filtered samples and compares them to a prefixed energy threshold. If the energy exceeds this threshold, the system initiates a thud component and the glass-breakage-detection algorithm. The digital lowpass filter must be small yet effective, so the sampling frequency for these initial samples remains at only 4 kHz. However, this section of the algorithm uses an antialiasing filter with a cutoff frequency of 2 kHz rather than 20 kHz for typical antialiasing filters.

The glass-breakage-detection algorithm is more complex than the thud detection and includes two signal-analysis parts. One is the first stage of processing and occurs on every sample once the algorithm detects a thud. This stage uses a 20-kHz antialiasing filter and increases the ADC's sampling frequency to 40 kHz. The operations for this stage include signal averaging, zero-crossing detection, and peak detection, which occur for approximately 60 msec, or approximately 2400 samples. Once the first stage is complete, the second stage initiates to complete the entire signal analysis.

Figure 5 shows a signal representa-

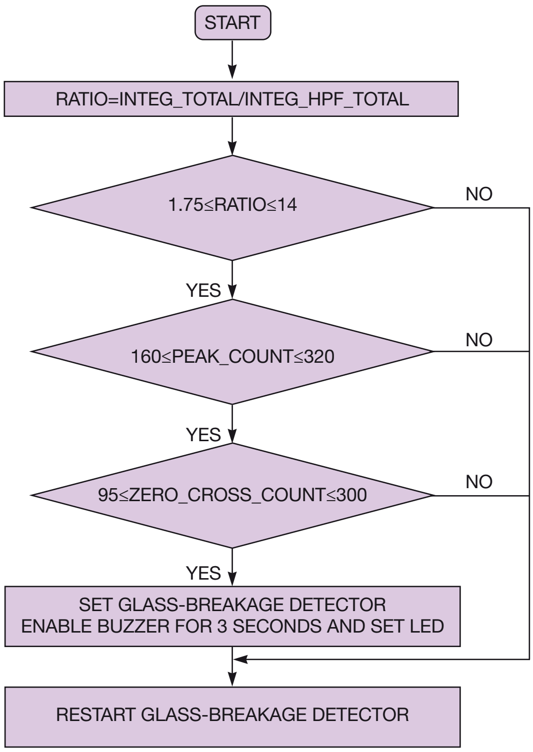

Figure 7 Once the first stage is complete, the second stage initiates to complete the entire signal analysis.

tion during the first stage, and **Figure 6** shows the software flow. The algorithm's p(n) signal denotes the incoming samples, which pass through a simple moving-average filter to reduce noise, yielding the s(n) signal. Integration of the p(n) signal uses only positive samples to calculate signal energy, integ_total, for use in the second processing stage. The s(n) signal then receives peak and zero-crossing counts. To extract the high-frequency components of the incoming signal, you use a highpass filter with a cutoff frequency of one-fourth the sampling frequency and subject each sample of p(n) to this filtering. Simultaneously, only the positive samples of the filtered output are accumulate in the result, integ_HPF_total, which the second stage uses. Stage 1 filtering occurs on every sample and must be complete before the arrival of the next sample, p(n+1), for real-time operation, implying that the total amount of CPU cycles available is only the CPU frequency divided by 40 kHz. Filtering is generally a time-consuming operation. To achieve efficiency, both the low-

THE ALGORITHM COMPUTES THE RATIO OF TOTAL SIGNAL ENERGY TO HIGHPASS-FILTERED ENERGY AND CHECKS THE RESULTS AGAINST A THRESHOLD.

pass filter and the highpass filter in Stage 1 use lattice-wave digital filters and Horner's algorithm for thud detection.

Once the first stage of signal analysis processes 60 msec of data, the algorithm proceeds to the second stage of processing. The second stage does not require real-time operation (**Figure 7**). The end of the second stage of signal analysis confirms whether any glass breakage has actually occurred.

The algorithm computes the ratio of total signal energy to highpass-filtered signal energy and checks the results against a threshold. Results show a ratio of 1.75-to-14 for a number of glass-breakage sounds. Similarly, the algorithm checks the number of peaks if it is between 160 and 320 and whether the number of zero crossings is between 95 and 300. A valid glass breakage occurs if the results satisfy each of these conditions. If even one of the conditions fails, the glass-breakage detector reinitializes and returns to activity detection. You

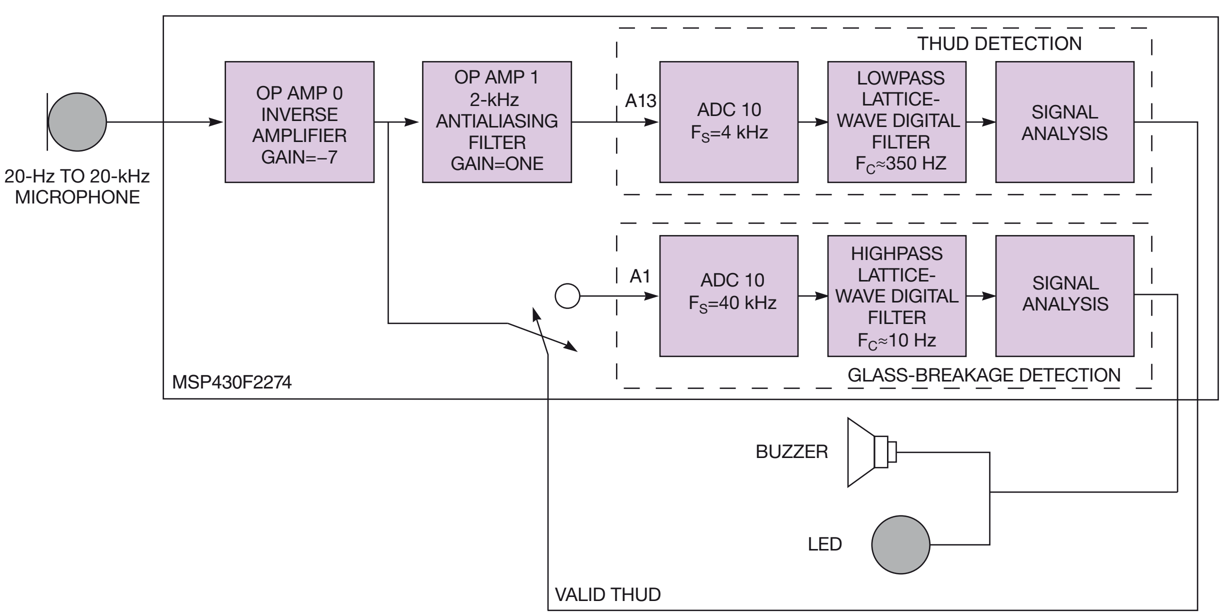

Figure 8 Texas Instruments' low-power, 16-bit MSP430F2274 microcontroller operates at frequencies as high as 16 MHz.

must slightly tweak these thresholds and ranges depending on the room's acoustics, the detector's location, noise in the environment, and other factors.

MICROCONTROLLER IMPLEMENTATION

The ultralow-power MSP430-microcontroller platform from Texas Instruments (www.ti.com) comprises a variety of devices, including the 16-bit MSP430F2274 microcontroller, which operates at frequencies as high as 16 MHz (**Figure 8**). It also has an internal low-power, low-frequency oscillator, which operates at 12 kHz at room temperature; two 16-bit timers; and a 10-bit ADC 10, which supports conversion rates as high as 200 kHz. The ADC works with on-chip, software-configurable operational amplifiers 0 and 1 for analog-signal conditioning. The device consumes 0.7 μA of current during standby-mode and 250 μA during active mode, making it a good choice for battery-powered applications. Because the chosen microphone has a passband of 20 Hz to 20 kHz and the MSP430F2274 integrates only two op amps, you could remove the 20-kHz antialiasing filter from the implementation. Although this removal violates sampling theory, the results do not vary with this absence. However, if another op amp is available, the filter can still be part of this setup.

Op amp 0 works as an inverting amplifier with a gain of seven to provide amplification to the microphone output. Op amp 1 works as a unity-gain lowpass filter, which is a second-order Butterworth type using the Sallen-Key architecture. The filter has a 3-dB cutoff frequency at 2 kHz. The outputs of the two op amps internally connect to channels A_1 and A_{13}, respectively, of the MSP430.

CURRENT CONSUMPTION

The glass-breakage detector's current consumption depends on the low-power modes it uses during its operation and selective turn-on and turn-off of its peripherals. The current-consumption profile of the implementation on the MSP430 for the three modes of operation is discussed in the online version of this article at www.edn.com/ms4375. **Table 1** provides a list of the peripherals and clocks that are on during various modes of operation. These peripherals are significant contributors to the overall current consumption during each stage.

Two AAA batteries providing 800 mAhr of energy power the glass-breakage-detection board. Although it is difficult to predict the battery life of such an application, assuming no glass breakage, the total current consumption is approximately 80 μA to give a battery life of about 416 days. You can further increase battery life by increasing the wake-up for activity detection to more than 2.5 msec; however, this approach increases the possibility of missing a sound event. EDN

AUTHOR'S BIOGRAPHY

Kripasagar Venkat is a system-applications engineer for Texas Instruments (Dallas). He holds a master's degree from the University of Texas—Dallas with an emphasis in signal processing and communication.

16-Bit, 1Msps SAR ADC with 94dB SNR

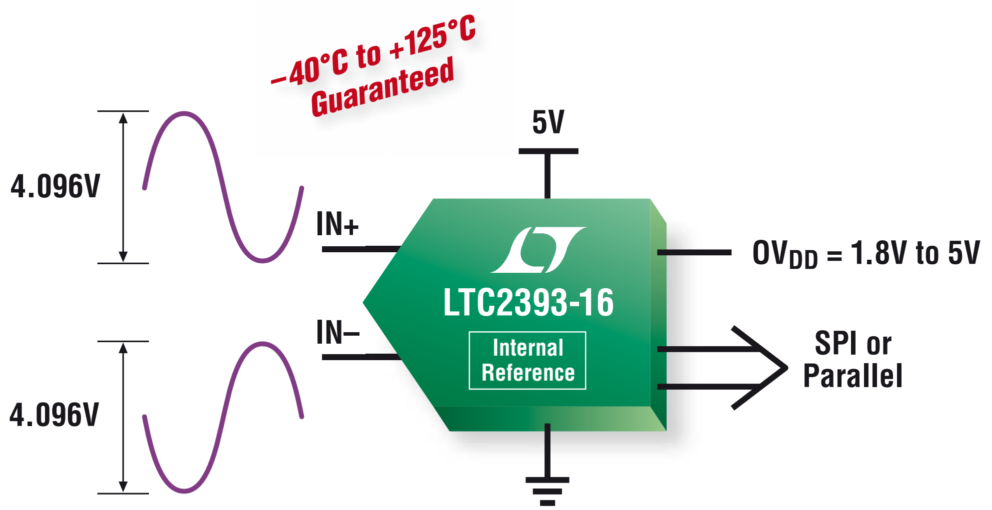

±4V$_{P-P}$ Differential Input Span ADC Excels in Noisy Environments

The LTC®2393-16's wide differential input range of ±4.096V and outstanding 94dB dynamic range ensures sufficient noise margin in hostile industrial environments. Operating from a single 5V analog supply, this ADC achieves 16-bit no missing code resolution at 1Msps with no latency. The LTC2393-16 includes an internal 20ppm/°C reference and is ideal for medical instruments, spectrum analysis, ATE, instrumentation and high-resolution data acquisition.

Features

- Pin- and Software-Compatible Family:
  - LTC2393-16: 1Msps
  - LTC2392-16: 500ksps
  - LTC2391-16: 250ksps
- Single 5V Analog Supply, 1.8V to 5V Digital Supply
- ±2LSB INL (Max), ±1LSB DNL (Max)
- 20ppm/°C (Max) Internal Reference
- –40°C to +125°C Guaranteed Temperature Range
- Parallel & Serial Interfaces
- 48-Pin 7mm x 7mm QFN & LQFP Packages

16k Point FFT f_S = 1Msps, f_{IN} = 20kHz

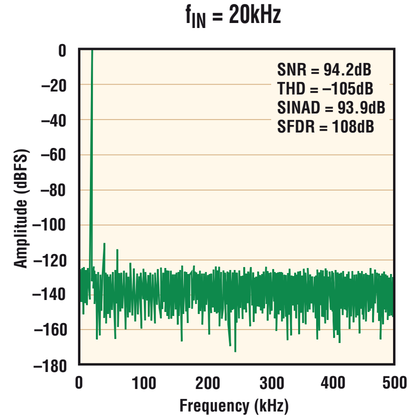

Info & Free Samples

www.linear.com/2393

1-800-4-LINEAR

Free Industrial Signal Chain Brochure

www.linear.com/indsolutions

LINEAR TECHNOLOGY

designideas

READERS SOLVE DESIGN PROBLEMS

Current monitor compensates for errors

Chau Tran and Paul Mullins, Analog Devices, Wilmington, MA

You sometimes need to measure load currents as large as 5A in the presence of a common-mode voltage as high as 500V. To do so, you can use Analog Devices' (www.analog.com) AD8212 high-voltage current-shunt monitor to measure the voltage across a shunt resistor. You can use this circuit in high-current solenoid or motor-control applications. **Figure 1** shows the circuit, which uses an external resistor and a PNP transistor to convert the AD8212's output current into a ground-referenced output voltage proportional to the IC's differential input voltage. The PNP transistor handles most of the supply voltage, extending the common-mode-voltage range to several hundred volts.

An external resistor, R_{BIAS}, safely limits the circuit voltage to a small fraction of the supply voltage. The internal bias circuit and 5V regulator provide an output voltage that's stable over the operating temperature range, yet it minimizes the required number of external components. Base-current compensation lets you use a low-cost PNP pass transistor, recycling its base current, I_B, and mirroring it back into the signal path to maintain system precision. The common-emitter breakdown voltage of this PNP transistor becomes the operating common-mode range of the circuit.

The internal regulator sets the voltage on COM to 5V below the power-supply voltage, so the supply voltage for the measurement circuit is also 5V. Choose a value for the bias resistor, R_{BIAS}, to allow enough current to flow to turn on and continue the operation of the regulator.

▶What are your design problems and solutions? Publish them here and receive $150! Send your Design Ideas to edndesignideas@cancom.com.

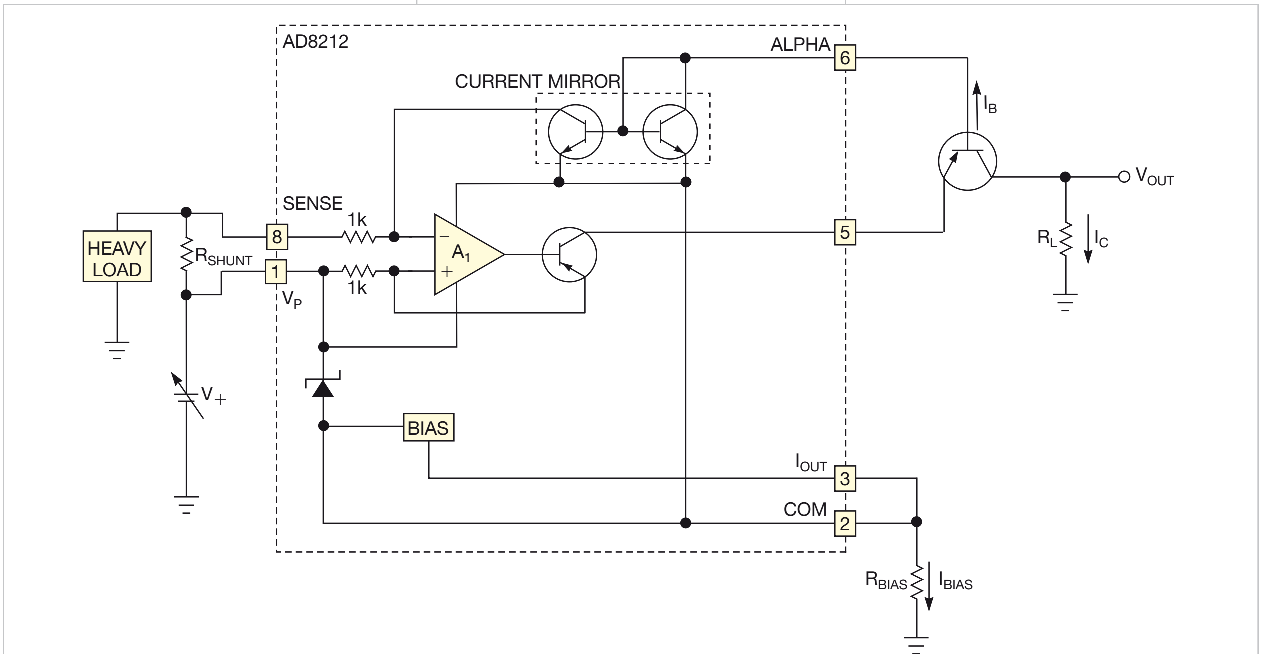

Figure 1 An external PNP transistor lets you operate the circuit at high voltages.

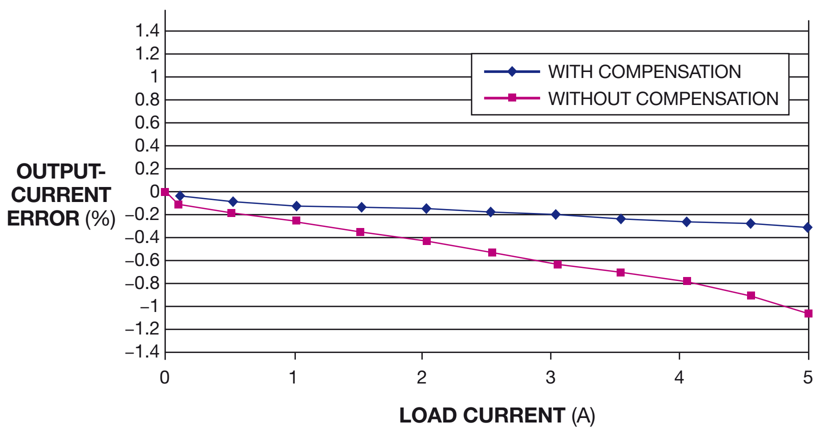

Figure 2 Internal base-current compensation reduces error.

For high-voltage operation, set I_{BIAS} at 200 µA to 1 mA. The low end ensures the turn-on of the bias circuit; the high end is limited, depending on the device you use.

With a 500V battery and an R_{BIAS} value of 1000 kΩ, for example, $I_{BIAS}=(V_+-5V)/R_{BIAS}=495V/1000\ k\Omega=495\ \mu A$.

The circuit creates a voltage on the output current approximately equal to the voltage on COM plus two times the V_{BE} (base-to-emitter voltage), or $V_+-5V+2V_{BE}$. The external PNP transistor withstands two times the base-to-emitter voltage of more than 495V, and all the internal transistors withstand voltages of less than 5V, well below their breakdown capability.

Current loss through the base of the PNP transistor reduces the output current of the AD8212 to form the collector current, I_C. This reduction leads to an error in the output voltage. You can use a FET in place of the PNP transistor, eliminating the base-current error but increasing the cost. This circuit uses base-current compensation, allowing use of a low-cost PNP transistor and maintaining circuit accuracy. In this case, current-mirror transistors, the AD8212's internal resistors, and amplifier A_1 combine to recycle the base current.

Figure 2 shows a plot of output-current error versus load current with and without the base-current-compensation circuit. Using the compensation circuit reduces the total error from 1 to 0.4%. You should choose the gain of the load resistor, R_L, to match the input voltage range of an ADC. With a 500-mV maximum differential-input voltage, the maximum output current would be 500 µA. With a load resistance of 10 kΩ, the ADC would see a maximum output voltage of 5V.EDN

Buck regulator handles light loads

Justin Larson and Frank Kolanko, On Semiconductor, East Greenwich, RI

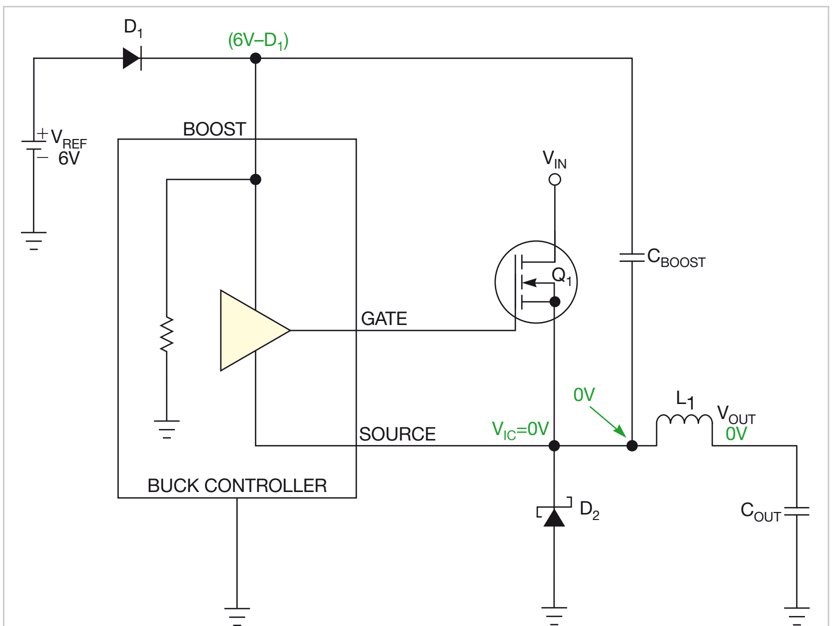

Figure 1 Many buck-controller ICs use a floating-gate driver.

Buck regulators operating in CCM (continuous-conduction mode) have straightforward operation, allowing for easy calculation of output voltage and system design. However, lightly loaded buck regulators operate in DCM (discontinuous-conduction mode), and their operation is more complicated. The duty cycle changes from a ratio of the output voltage and the input voltage. A regulator that reduces a 12V input to 6V has a 50% duty cycle. When the regulator is too lightly loaded to keep some current continuously flowing in the inductor, it enters DCM. The duty cycle changes to a complex function of inductor value, input voltage, switching frequency, and output current, greatly slowing the control-loop response.

Many buck-controller ICs use a floating-gate driver (**Figure 1**). You use a separate supply reference voltage, V_{REF}, for high efficiency. During start-up, it powers the NFET gate driver to a diode drop below the reference voltage. Sufficient voltage is available to drive the

gate of the FET because the initial conditions dictate 0V on the output and on the source of FET Q_1.

During CCM, current always flows through the inductor. Q_1 or D_2 supplies this current during the flyback event that Q_1's turn-off causes (**Figure 2**). The flyback event creates a voltage at the source of Q_1, and the drop across D_2 limits this voltage, making it a negative voltage with respect to ground. Sufficient voltage is available to drive Q_1 because the C_{BOOST} capacitor boosts the gate voltage. This boost provides a high voltage to the boost pin and the resultant negative voltage on the Q_1 source.

The system enters DCM when the load drops to the point at which the average current demand is less than one-half the current ripple. Diode D_2 prevents reverse current in the inductor. Depending on the chip you use, the output may overshoot due to the slower response time of the control loop. The regulator may also miss pulses and generally operate unpredictably. After Q_1 turns off, C_{BOOST} starts to bleed down through the boost pin and D_1 (**Figure 3**). The extended off time of Q_1 in DCM starts to discharge the C_{BOOST} capacitor. At approximately 3V across C_{BOOST}, Q_1 does not turn on until the output capacitor, C_{OUT}, discharges adequately to provide a lower voltage on the source of Q_1 than that of the boost pin through D_1. This behavior is unacceptable in a voltage regulator.

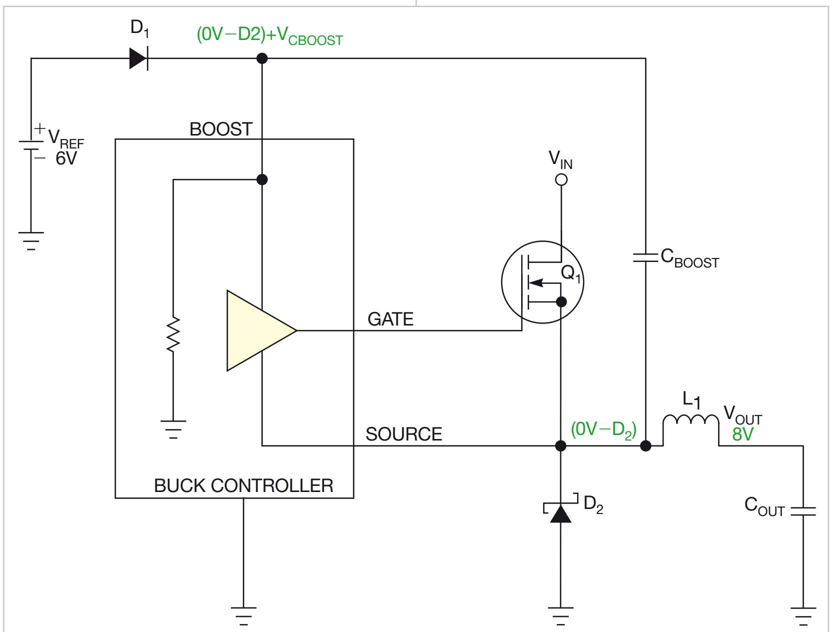

Figure 2 During CCM, current always flows through the inductor. Q_1 or D_2 supplies this current during the flyback event that Q_1's turn-off causes.

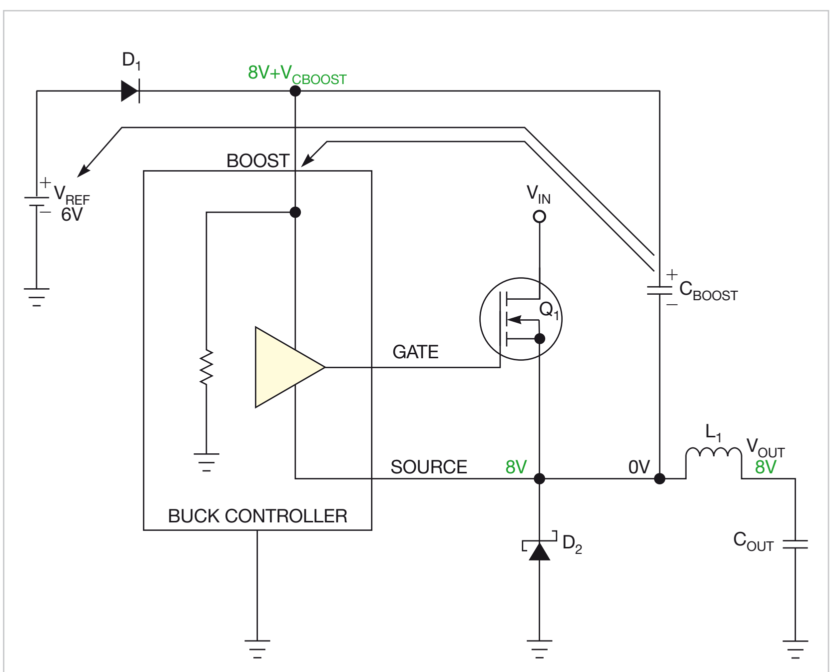

Figure 3 The C_{BOOST} capacitor discharges when the regulator goes into DCM.

YOU DON'T KNOW THE TEMPERATURE COEFFICIENT OF THE CURRENT INTO THE BOOST PIN, SO YOU SHOULD ALSO CHECK OPERATION AT LOW TEMPEATURE.

High temperatures create a situation with higher leakage currents. You don't know the temperature coefficient of the current into the boost pin, so you should also check operation at low temperature. Evaluate the system to determine the lowest capacitor value, using this result in your worst-case evaluation simulations. You can thus ensure that the design will operate in DCM by increasing the value of C_{BOOST}. You could also increase the reference voltage to which D_1 connects. You may want to consider replacing D_1 with a low-leakage Schottky diode. If none of these approaches results in reliable operation, you can switch to an IC that uses a gate driver referenced to ground or modify your design to use a synchronous-buck architecture. EDN

Sense multiple pushbuttons using only two wires

Bernhard Linke, Maxim Integrated Products Inc, Dallas, TX

Keyboards and numeric keypads often provide the user interface for electronic equipment, but many applications require only a few pushbuttons. For those applications, you can monitor multiple pushbuttons over a single pair of wires (**Figure 1**).

The multichannel 1-Wire addressable switch, IC_1, provides PIO (input/output ports) P0 through P7, which in this application serve as inputs. The 1-MΩ R_{PD} resistors connect these ports to ground to ensure a defined logic-zero state. Diode/capacitor combination D_1/C_1 forms a local power supply that steals energy from the 1-Wire communication line. Pressing a pushbutton connects the corresponding port to the local supply voltage, which is equivalent to logic one. This change of state sets the port's activity latch (**Reference 1**).

As a 1-Wire slave device, IC_1 doesn't initiate communications. Instead, the master device—typically, a microcontroller—polls the 1-Wire line. To minimize overhead, IC_1 supports conditional search, a 1-Wire network function. Before using that function, however, you must configure IC_1 according to the needs of the application. That configuration includes channel selection, which defines the qualifying input ports; channel-polarity selection, which specifies the polarity of the qualifying ports; choosing a pin or an activity latch for the port; and specifying whether the device will respond to activity at a single port, an OR, or at all ports, an AND.

Consider, for example, that IC_1 shall respond to a conditional search if it detects activity at any of the eight ports. This search requires a channel-selection mask of 11111111b for address 008Bh. The numeral one indicates that IC_1 has selected a port. This search also requires a channel-polarity selection of 11111111b for address 008Ch, where the numeral one indicates a high level, and a control/status register setting of 00000001b for address 008Dh, which selects the port's activity latch as a source and specifies OR as the conditional search term—that is, activity at any port.

After power-up, the configuration data must be loaded into IC_1 using the write-conditional-search-register command. Next, the channel-access-write command, with FFh as a PIO output-data byte, defines the ports as inputs. Subsequently, the issue of a reset-activity-latches command completes the configuration. IC_1 is now ready to handle pushbutton activity.

After you configure IC_1, the application software enters an endless loop, in which a conditional-search command follows a 1-Wire reset. With no pushbutton activity, IC_1 does not respond, as

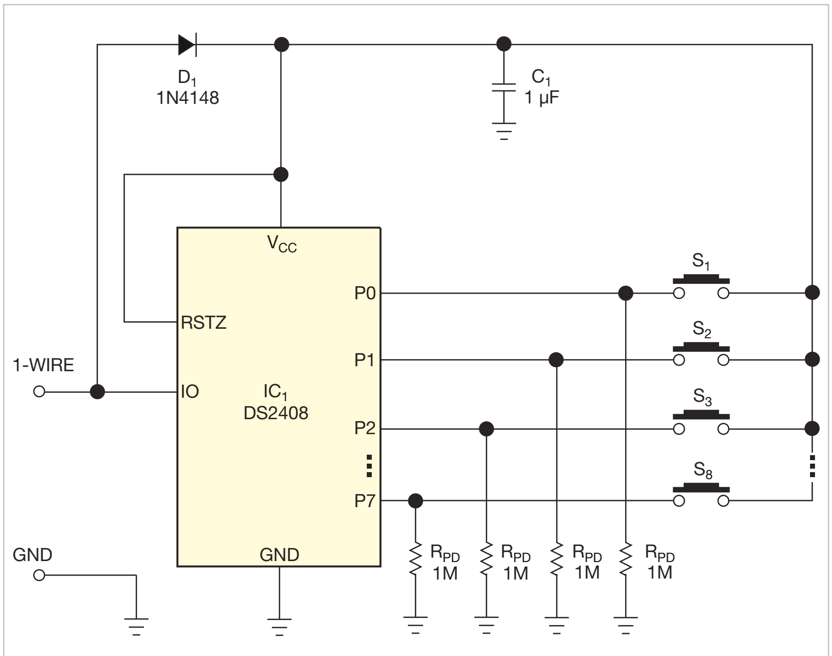

Figure 1 This circuit connects to a microcontroller and can monitor eight pushbuttons using only two wires.

a logic one indicates, for the 2 bits immediately after the conditional-search-command code. In that case, the microcontroller cancels the conditional search and starts over.

If IC_1 responds to the conditional search, the first 2 bits will be one and zero, representing the least-significant bit of the device's family code, 29h, in its true and inverted forms. In that case, the microcontroller should complete the conditional-search flow, which comprises a 192-bit sequence. Next, the microcontroller reads from IC_1 by issuing a read-PIO-registers command using 008Ah, the address of the PIO-activity-latch-state register. The microcontroller then issues a 1-Wire reset, a resume command, and a resume-and-reset-activity-latches command. It then returns to the endless loop, polling for the next pushbutton event.

If IC_1 responds and no other 1-Wire slave is connected, the microcontroller could cancel the conditional search after reading the first 2 bits, issue a 1-Wire reset, issue a skip-ROM command, and then read the PIO-activity-latch-state register. Next, it must issue a 1-Wire reset, a skip-ROM command, and a reset-activity-latches command before returning to the endless loop.

The code read from the PIO-activity-latch-state register tells which button was pressed. If you press S_1, the data is 00000001b; if you press S_2, it is 00000010b; and so forth. At least one of the 8 bits will be one. If you press several buttons after the last reset-activity-latches command, several bits are one. The application software must then decide whether such a condition is valid. In the simplest case, one-of-eight code, the software considers all codes that have several bits at one as invalid.

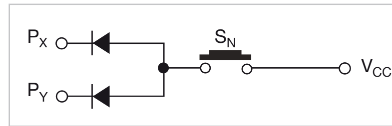

Figure 2 This circuit can monitor 28 additional pushbuttons if you use diodes to connect them to two ports.

You can expand this concept to more than eight pushbuttons. Instead of associating one pushbutton with one port, you can associate additional pushbuttons with two simultaneously activated ports, representing two-of-eight code (**Figure 2**). If another pushbutton activates P_X or P_Y, the diodes prevent that activity from propagating to other ports. Again, the application software must check the code it reads from the PIO-activity-latch-state register to decide whether it is valid. The theoretical limit of this concept is 255 pushbuttons, which require combinations of two, three, four, five, six, seven, or eight diodes per additional pushbutton. When the cost of diodes for each additional pushbutton begins to exceed the benefits, you will find that adding another DS2408 is more cost-effective. EDN

REFERENCE

1 "DS2408 1-Wire 8-Channel Addressable Switch," Maxim Integrated Products Inc, www.maxim-ic.com/ds2408.

Tricolor LED emits light of any color or hue

Marián Štofka, Slovak University of Technology, Bratislava, Slovakia

The human eye can see any color as a mixture of blue, red, and green. The circuit in **Figure 1** produces all three colors through an Avago (www.avagotech.com) ASMT-YTB0 tricolor LED. You can produce a wide range of colors by varying the current in the blue, red, and green LEDs.

The collector outputs of bipolar differential stages form the current sources. A classic symmetrical differential stage with two equal bipolar transistors is a backbone of almost all bipolar analog ICs. In this case, however, the differential stage is asymmetrical, with a 2-to-1 collector-current distribution instead of the common 1-to-1 ratio at 0V base-voltage difference. The circuit produces the 2-to-1 current ratio by paralleling a third equal transistor, Q_3, to Q_1. The common collector of the paralleled transistor pair connects to the common emitter of the Q_4/Q_5 differential stage. Thus, the base differential voltages equal 0V at both the stages, and collector currents I_R, I_G, and I_B are almost equal.

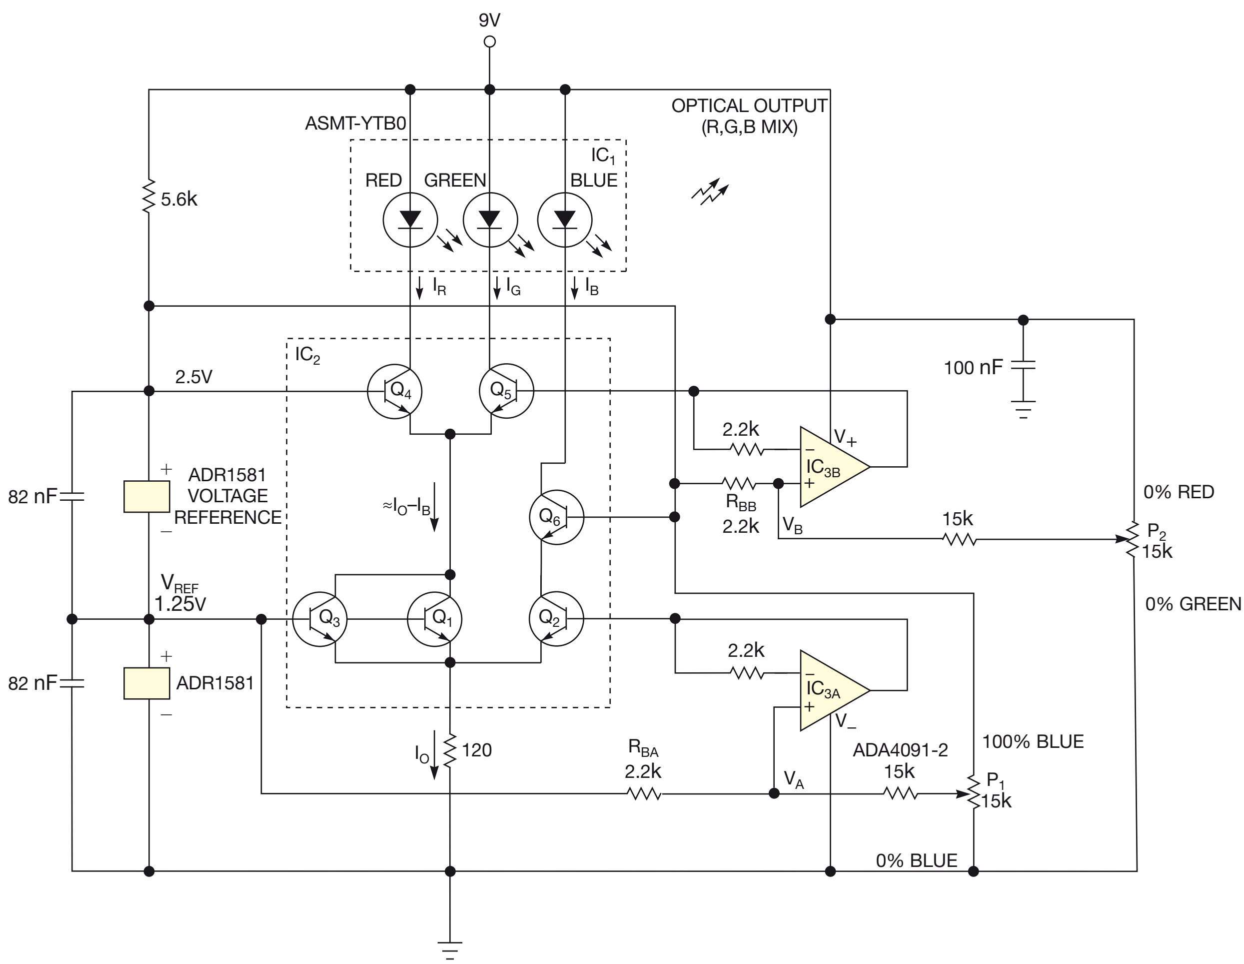

Figure 1 Potentiometers P_1 and P_2 let you control the color of emitted light.

The differential stages let you vary I_R, I_G, and I_B over a range of 0 to I_O, where $I_R+I_G+I_B \approx I_O$=4.43 mA. This value is approximate because $I_R+I_G+I_B$ is lower by a relative value of 3/β, where β is a current gain of the bipolar transistors. The relative error is less than 1%. Transistor Q_6 equalizes Q_2's collector voltage with those of the Q_1 and Q_3 collectors. This approach preserves the matching of the base-emitter voltages of Q_1, Q_2, and Q_3. The base currents of bipolar transistors in this case can reach to as much as 100 µA. For this reason, you route the color and hue control voltages, V_A and V_B, which you derive from resistive potentiometers P_1 and P_2, to the bases of Q_2 and Q_5 through voltage-follower-connected op amps IC_{3A} and IC_{3B}, two halves of an Analog Devices' (www.analog.com) ADA4091-2. The ADA4091-2 has low power consumption and input offset voltage of less than 500 µV with a typical value of 80 µV.

The ADA4091-2 has a maximum input bias current of 65 nA, which causes a negligible voltage drop on resistors R_{BA} and R_{BB}. This voltage drop is less than 130 µV. You can achieve even more accuracy by inserting resistors of the same value as R_{BA} between the respective inverting inputs and outputs of both the A and the B followers. This step brings reduction of input-bias-current-caused errors to one-sixth worst case—down to 1/600.

Potentiometer P_1 controls the blue LED's intensity. At the upper-end position, when the LED is 100% blue, transistors Q_2 and Q_3 are off, which turns off Q_4 and Q_5. Thus I_O flows solely through Q_2 and Q_6. The red and green LEDs are therefore off. When P_1's wiper is at 0V, output current flows exclusively through paralleled Q_1 and Q_3 and distributes itself to Q_4 and Q_5, depending on the position of the wiper of potentiometer P_2. With P_2's wiper at its upper end, the circuit emits 100% green light. At 0V, the emitted light is fully red. An intermediate position of the wiper yields a mixture of red and green. By moving P_1's wiper from the ground position, the circuit produces a mixture of red, green, and blue.

Transistors Q_1, Q_2, and Q_3 should tightly match. You need a difference in base-emitter voltages of less than 1.5 mV. The same requirement holds true for the Q_4/Q_5 pair. Matching requirements are less stringent for Q_6. You should use a bipolar NPN matched-transistor pair for Q_1 through Q_6, or at least Q_1 through Q_5, whereas Q_6 is a single transistor. Eventually, you can use three matched-transistor pairs. EDN

High-reliability-component market maintains growth

The market for high-reliability components was one of the few growth sectors during the recent economic downturn. The concentration of military, medical, space, and aerospace, while not recession-proof, experienced mostly moderate growth over the past two years. Although much of this market involves serving legacy systems, new advances in satellite technology, medical systems, and down-hole drilling are breathing new life and new products into the high-reliability-component market.

Much of the high-reliability market is still anchored in supporting these legacy systems. "We started off with a focus almost exclusively on the high-reliability market in 1971," says Michael Knight, vice president of product and supplier marketing at TTI Inc (www.ttiinc.com). "Now, these components make up about 20% of our overall business."

Texas Instruments Inc (www.ti.com) started 30 years ago to produce high-reliability components. "Most of our customers were in defense and aerospace," says Brad Little, product-line manager for TI's high-reliability team. "In the last few years, our customer base has become much broader." Newer high-reliability applications, Little says, include oil exploration, down-hole drilling, undersea cabling, and railway lighting, all of which require product longevity in high-temperature environments. The

medical market also requires high-reliability devices for use in sterilization equipment or imaging and radiation applications.

These applications have helped the high-reliability market move beyond simply providing legacy components. "There are on average 100 satellites launched each year, and these are long-term, two- to five-year projects," says Little. All satellites require replacement over time, and today's economy relies strongly on satellite-data capability.

One of the biggest challenges with the space environment is the bombardment of potentially damaging radiation. "We're introducing a lot of new radiation-hardened devices for space, medical, and harsh environments," says Little. Eliminating excess hydrogen in the IC-manufacturing process—a costly method—can improve radiation tolerance; reduce leakage, thus allowing a device to operate over higher temperatures; and improve reliability from a 10-year expected lifetime to a 20-year operational lifetime, according to Little.

Hardening components for harsh radiation environments can drastically affect the price of components. "A lot of satellites go up every year, so a lot of radiation-hardened parts are coming online," says Bill Toumey, supplier and programs manager at the aerospace and defense group at Arrow Electronics Inc (www.arrow.com). The commercial, high-reliability, and radiation-hardened versions of these parts sell for \$5, \$15, and \$500, respectively, even though the versions offer the same functions.

The satellite market is keeping a number of component manufacturers interested in developing new high-reliability products. "The satellite business is going to grow significantly, so we're concentrating on building for hard environments, wide temperature ranges, and radiation resistance," says Ron Demcko, applications-engineering manager at AVX Corp (www.avx.com), a component manufacturer. "We're concentrating on this area because we know it's going to be steady," he adds. "We're changing the configuration of parts, and we're changing material systems to increase the performance of our components."

The market for high-reliability components is generally resistant to economic ups and downs, as opposed to the volatile computer and consumer-electronics markets. Whereas component prices and production went through wild swings over the past three years, high-reliability components did not. "Military semiconductors stayed robust and continued to grow even through the recession," says Bryan Brady (**photo**), vice president of defense and aerospace at Avnet Inc (www.avnet.com). He explains that funding for legacy programs still using military-specification parts drove much of that growth and that prices for the parts have remained stable.

Finding leaded components continues to be a challenge for those buying high-reliability parts. Although many component manufacturers are supporting legacy parts in leaded versions, few have created leaded versions of their new components. "By definition, high-reliability parts are leaded," says Brady. "Over the past 10 to 15 years, the commercial plastics and COTS [commercial-off-the-shelf] product has become dominant in the military." However, new designs are almost all commercial products, and many of those new parts do not contain lead, he notes. "Sometimes, the military is using mitigation strategies, such as solder dipping and reballing, to guard the component against tin whiskers and other solder-reliability concerns."

A number of component manufacturers have created leaded versions of their new components. "Any part we produce commercially will also be introduced with tin-lead alternatives," says Chris Reynolds, technical-marketing manager at AVX.—**by Rob Spiegel**

Introducing EDN's New Microsite on LEDs

Your complete source for up-to-the-minute information on LEDs and their applications.

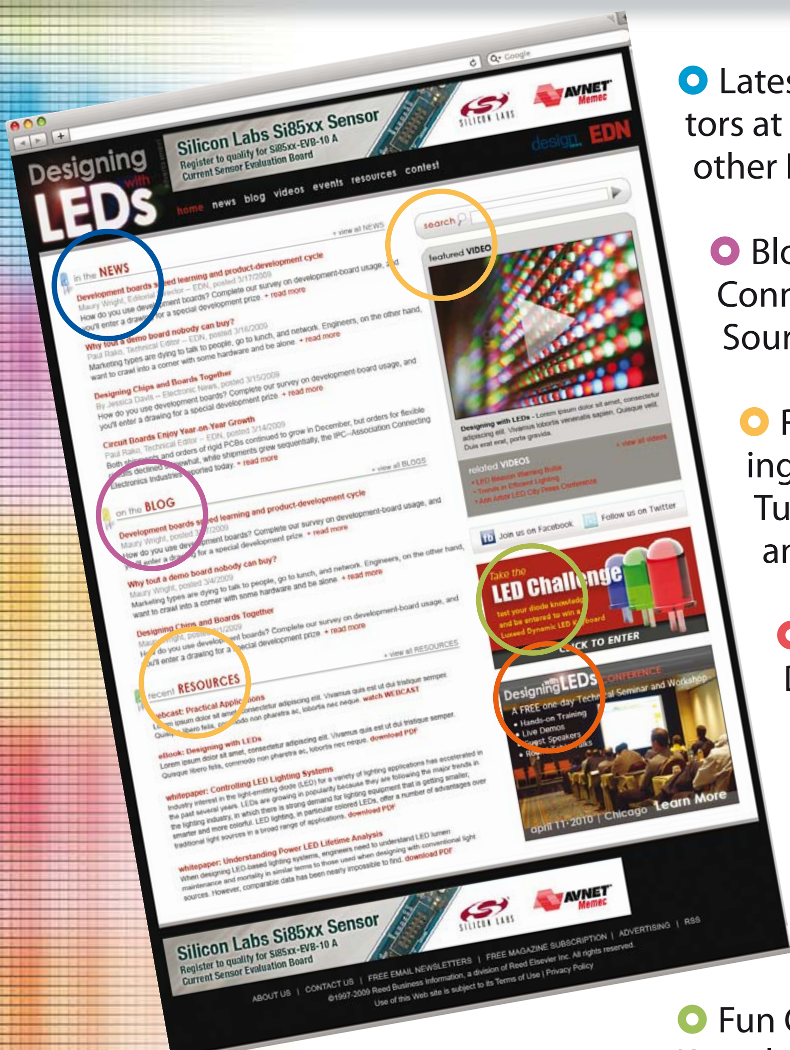

- Latest News from the Editors at EDN, Design News and other Engineering Magazines
- Blogs, including Margery Conner's Popular "Power-Source" Blog
- Resource Center, Including White Papers, eBooks, Tutorials, Webcasts and More
- Videos, Including Design Tips, Application Demos, and More
- Updates on Future In-Person Workshops and Seminars, like this one!
- Fun Challenges to Test Your Knowledge of LEDs

Check It Out Today at www.edn.com/leds

productroundup

SWITCHES AND RELAYS

Rocker switches come in illuminated and two-tone versions

The CG and CL series of rocker switches target use in household appliances, instrument panels, industrial controls, computers, and peripherals. The UL- and CSA-approved single-pole switches feature two-tone and multicolored illuminated actuators. The CG series devices come with two-tone actuators in five color combinations or with red-, amber-, or green-illuminated actuators. Each switch has an electrical life of 10,000 make-and-break cycles at full load. The CL series is available with illuminated actuators in red, amber, and green options. The contact ratings for both series are 16A at 125V ac or 10A at 250V ac, and operating-temperature range is –20 to +85°C for the CG series and –20 to +80°C for the CL series.

C&K Components, www.ck-components.com

Load switches ease power-design challenges

The FPF1038 and FPF1039 load-management switches provide a one-chip, monolithic approach to the challenge of reducing inrush current when the switch disconnects loads with more than 100-μF output capacitances. Both switches integrate a slew-rate-controlled MOSFET switch with 21-mΩ typical impedance. Targeting use in embedded-system applications, the switches also feature power consumption of less than 1 μA, a low-discharge path, ESD protection, and GPIO/CMOS-compatible-enable circuitry. The devices also feature a 1.2 to 5.5V input-voltage operating range, which aligns with supply rails for the embedded processors, custom ASICs, and FPGAs these applications use. The optimized slew-rate-controlled turn-on characteristics with recovery time of 2.7 msec prevent voltage droop on supply rails with bulk capacitances as large as 200 μF. The switches cost 60 cents (1000) each.

Fairchild Semiconductor, www.fairchildsemi.com

Long-distance reflective switch has adjustable sensitivity

The OPB725A-18Z reflective switch uses an infrared LED and a logic-output sensor and features a reflective-sensing distance as long as 24 in., depending on the object. The switch ensures detection of an 8×8-in., 90%-diffuse reflective white card at a

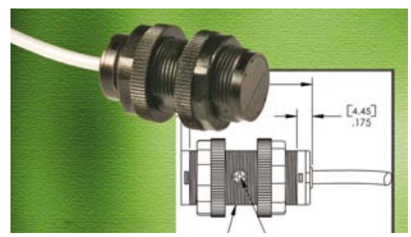

distance of 18 in. A trimming potentiometer, accessible through a 4.3-mm-diameter hole in the sensor's body, sets the sensitivity to maximum using clockwise rotation. The device features an 850-nm LED and comes with 48-in. (122-cm) #26 AWG with mounting nuts. Power dissipation is 250 mW, maximum collector voltage is 30V, and collector dc current is 50 mA. Operating temperature ranges from 0 to 50°C, and the price is $15.44 (5000).

Optek Technology, www.optekinc.com

Tamper-resistant switches come in illuminated and nonilluminated versions

Devices in the LP01 series of short-body, illuminated and nonilluminated, tamper-resistant push-button switches have a behind-panel depth of 14 mm. The cap diameter is 13.6 mm, making the overall diameter, including housing, just 20 mm. The switches come fully assembled, and the illuminated versions feature red-, green-, or amber-LED illumination and a translucent white cap. With a mini-

mum mechanical life of 500,000 operations, these devices target use in high-usage applications. They have ratings of 3A at 125/250V ac or 3A at 30V dc. Prices start at $3.90 (2500) each for the nonilluminated version and $4.81 for the illuminated version.

NKK Switches, www.nkkswitches.com

EDN ADVERTISER INDEX

EDN provides this index as an additional service. The publisher assumes no liability for errors or omissions.

EDN productmart

This advertising is for new and current products.

The devil's in the details

DANIEL VASCONCELLOS

About 10 years ago, while working as a design engineer for a capacitor-manufacturing company, I was responsible for designing new products that were not in the company's product line. One day, my boss passed me the specs for a new MPP (metallized-polypropylene) shunt-capacitor form factor that we needed to develop. The specs had all the relevant details on size, capacitance, and dissipation and form factor. On the first pass, however, I noticed that the details on the application for this device were missing. That is, we knew what the customer wanted but not the application—beyond circuit filtering—for the parts.

Nevertheless, we established our priorities and set out to design, fabricate, test, and qualify a sample set before the end of the week. The customer urgently required delivery because evaluation was part of the company's business process. We also knew that the design team at the customer's end would surely know what the application required, and this part was only a small component in a big system. So we proceeded to fulfill the customer requirements. We chose the standard design parameters and materials. The manufacturing team had to manufacture the sample set under close supervision; we then tested and qualified the samples and agreed that they were ready for shipment.

Nearly a month passed, and, as often happened, our marketing team could elicit no feedback on the samples from the customer. They were probably collecting dust on someone's desk until they were actually needed. Surprisingly, however, nearly two months after we sent them out, we received eight of the 10 samples for field-failure analysis. The task of opening the burned and badly scarred parts and doing a critical failure analysis on them was daunting.

Nonetheless, we proceeded and determined that the failure had occurred in a polypropylene puncture that had not self-healed and had then gone on to cause the capacitor to burst. We could find no signs of any material or manufacturing defect. Besides, we had thoroughly conditioned the samples during burn-in and had shipped only the good units. We had also run samples at our in-house lab for 1000-hour testing, and we had not seen any failures.

The failures received a lot of attention from our company's top management because the customer was strategic to our business. It was now inevitable that our marketing team would have to bite the bullet and push the customer to share the application's details with us.

A week later, we received an e-mail from our company's field-applications team. The customer was using the capacitors to filter motor-supply noise for dc-motor applications. It immediately hit us: The motor spikes were causing heavy overvoltage impulses across the capacitor. These spikes would breach the breakdown-threshold voltage of polypropylene capacitors, and the capacitors' self-healing would go into overdrive, leading to the capacitors' bursting. Two simple solutions were available. We could either adequately derate the capacitors' voltage or overdesign the parts to meet the customer's application requirements. Either of these options would incur additional cost. Once we explained the failure and design criteria, the customer opted for us to implement both approaches.

The lessons we learned were, first, consider the application and environmental details of a device early in the design process and, second, customers themselves may fail to impart to you the small but essential requirements of their applications. Finally, in the corporate world, speed means money—usually at the cost of careful completion.EDN

Brian Fernandes is a senior design engineer in Singapore.

+ www.edn.com/tales

2 GHz Clock Generator

CG635...$2490 *(U.S. list)*

- ***Square wave clocks from DC to 2.05 GHz***
- ***Random jitter <1 ps (rms)***
- ***80 ps rise and fall times***
- ***16-digit frequency resolution***
- ***CMOS, LVDS, ECL, PECL, RS-485***
- ***Phase adjustment & time modulation***

Plot shows complementary clocks and PRBS (opt. 01) outputs at 622.08 Mb/s with LVDS levels. Traces have transition times of 80 ps and jitter less than 1 ps (rms).

The CG635 generates clock signals — flawlessly. The clock signals are fast, clean and accurate, and can be set to standard logic levels.

How fast? Frequency to 2.05 GHz with rise and fall times as short as 80 ps.

How clean? Jitter is less than 1 ps and phase noise is better than –90 dBc/Hz (100 Hz offset) at 622.08 MHz.

How accurate? Using the optional rubidium timebase, aging is better than 0.0005 ppm/year, and temperature stability is better than 0.0001 ppm.

You would expect an instrument this good to be expensive, but it isn't. You no longer have to buy an rf synthesizer to generate clock signals. The CG635 does the job better—at a fraction of the cost.

SRS

Stanford Research Systems

Phone: (408) 744-9040 · Fax: (408) 744-9049 · info@thinkSRS.com · **www.thinkSRS.com**